Artefactual Intelligence

Artefactual Intelligence

The development and use of cognitively congenial artefacts

David de Léon

Lund University Cognitive Studies 105

For further information, corrections and updates, see the author's web page at *http://www.lucs.lu.se/people/david.de.leon/*

Cover graphics: details from corrected teapot

Printed in Sweden by Grahns Tryckeri AB, Lund 2003.
ISBN 91-974741-1-8
ISSN 1101-8453
ISRN LUHFDA/HFKO-1013-SE
Lund University Cognitive Studies 105

For Ami and Boppa

… the world is peopled with objects. On its shores, their infinite throng, their collection appears to us, of course, rather indistinct and hazy.
But that is enough to reassure us. For, we feel it too, each one of them, at our mercy, in turn, can become our point of mooring, the limit on which we lean …

— Francis Ponge

Living in a material world
And I am a material girl
You know that we are living in a material world
And I am a material girl

— Madonna

CONTENTS

PUBLICATION HISTORIES

The publication histories for the papers included in the thesis are as follows:

Paper one

de Léon, D. (2003). Actions, artefacts and cognition: an ethnography of cooking. *Lund University Cognitive Studies, 104*.

Paper two

de Léon, D. (2002). Cognitive task transformations. *Cognitive Systems Research, 3(3),* pp. 449–459.

The abstract is also published as follows:

de Léon, D. (2001). Cognitive task transformations. In L. Chen and Y. Zhou (Eds.), *Proceedings of the Third International Conference on Cognitive Science* (p. 690). Beijing: Press of University of Science and Technology of China (USTC).

Paper three

de Léon, D. (2003). The cognitive biographies of things. *Lund University Cognitive Studies, 103*.

The paper is based on a presentation given at the *Doing Things with Things* workshop in Copenhagen Denmark in August 2001. It is presently published as a working paper at the department of Cognitive Science at Lund University and will eventually appear in a volume edited by Alan Costall and Ole Drier.

A popular version of the text, treating the same case study and many of the same issues, has been published in Swedish as:

de Léon, D. (2002). Att lägga tankarna på hyllan. In C. Fredriksson (Ed.), *Tingsligheter* (pp.12–21). Lund: Kulturens årsbok 2002.

Paper four

de Léon, D. (1999). Building thought into things. *Proceedings of the 3rd European Conference on Cognitive Science,* pp. 37–47.

Paper five

Hall, L., de Léon, D., & Johansson, P. (2002). The future of self-control: distributed motivation and computer-mediated extrospection. *Lund University Cognitive Studies, 95.*

The paper has also been submitted to *ACM Transactions on Computer–Human Interaction* (TOCHI), for a special issue on social issues and human–computer interaction.

ACKNOWLEDGEMENTS

A thesis such as the present one – which is composed of a collection of already published, or otherwise circulated, papers – has a fairly limited readership. Although I genuinely hope for one or two new readers, most of those who will hold this book in their hands will be colleagues, friends and family who have all previously trudged through the material reprinted here. I assume most will eschew a second reading of the papers and limit themselves to skimming through the introduction and flicking through the rest of the book looking at the pictures. I imagine, however, that all will read these acknowledgements.

The acknowledgements is often the only part of a thesis that I myself will read, usually in the tender hope of finding my own name and, if I'm lucky, briefly satisfying my sense of self-importance. On those occasions when I am indeed acknowledged in some way I experience a simple and direct pleasure: it is nice to know that one has made a difference. For me it is a pleasure that is enhanced by a childish delight in seeing my own name in print.

In what follows I hope to provide similar satisfaction for all of those who have helped, supported and encouraged me throughout my own project. There are many people to whom I am grateful, and without whom this thesis would simply not have been, or without whom writing it would have been so much less fun. All of you deserve to feel important and essential. I therefore sincerely hope that everyone will find themselves appropriately acknowledged.

*

Because they have contributed continuously and in innumerable ways from the very beginning I would like to start by thanking my advisor and my colleagues at the Department of Cognitive Science at Lund University. The climate at the department has been wonderful, as evidenced by the fact that most of my colleagues – past and present – have either been collaborators at some stage, or are now my close friends, or both.

My advisor, and head of the department, Peter Gärdenfors has created a dynamic and intellectually unfettered environment for which I am extremely glad to have been a part. Over time I have come to realise what a truly unusual and valuable place this is. I would like to thank Peter for this and for his sense of humour, broadmindedness and for permitting me to roam, sometimes pursuing what must have seemed rather strange and eccentric interests.

The weekly seminars held at the department have, more often than not, been productive and invigorating. Almost all my work has passed through the purgatory of our critical and hectic seminars at some point and I owe something to almost everyone who has participated in these over the years. There have been innumerable seminars and I therefore hope that I may be forgiven if someone should slip my mind. In addition to those who are singled out elsewhere in these acknowledgements let me take the opportunity to express my heartfelt thanks to: Christian Balkenius, Martin Bergling, Petra Björne, Ingar Brinck, Nils Dahlbäck, Pierre Gander, Agneta Gulz, Paul Hemeren, Måns Holgersson, Kenneth Holmqvist, Jana Holsánová, Nils Hulth, Per Johansson, Peter Kitzing, Lars Kopp, Maria Larsson, Jan Morén, Robert Palbo, Arne Svensk and Jordan Zlatev. You have all helped in one way or another.

Of all my colleagues in Lund none have helped or influenced me as much as my friend Henrik Gedenryd. Henrik was instrumental in helping me set course at the commencement of this project, feeding me key texts and discussing central issues pertaining to cognition and design. His close and critical reading of many of my papers has been invaluable and his own work on design and cognition inspirational. Sadly Henrik passed away as this project was nearing its completion. That he won't be around to read these acknowledgements (and to chastise me for many of the failings of this thesis) is both strange and sorrowful. I miss him very much, and I will miss the discussions that we were supposed to have over the years to come.

Annika Wallin is both friend and colleague and has been supportive from beginning to end. Annika and I have had almost daily discussions, ranging from the flippant to the profound. She has been an invaluable and meticulous discussion partner and has contributed greatly with her in-depth reading and ruthless comments of almost everything that I have written. I couldn't have wished for a better roommate and fellow traveller.

My dear friend, and also former colleague, Simon Winter has been a constant source of inspiration and encouragement. I have taken immense pleasure in our countless discussions and in Simon's creativity and way of thinking. I am also grateful for the intriguing, and often distracting, URL:s with which he has plied me over the years.

A daily source of joy are my friends, colleagues and sometime collaborators Lars Hall, Petter Johansson and Jens Månsson. In addition to an unending stream of playful lunches, and some very memorable and fruitful brainstorming sessions, I am thankful to them for allowing me to spring on them – usually unbid – various questions and half thoughts pertaining to my research. In this respect Lars Hall's encyclopaedic knowledge of cognitive science and adamant theoretical views have been particularly helpful, and Petter's rhetorical and strategical eye an invaluable resource. I would like to thank all three for what has been the most exciting, innovative and promising collaboration that I have ever been involved in. Powerful stuff can happen when we put our minds together.

Outside the department, my close friend Björn Nilsson has played an essential part in shaping my view of the relationship between people and things. Björn was instrumental at an early stage in alerting me to relevant work in anthropology and sociology, and our numerous and stimulating discussions about material culture and its relation to various facets of life have had a deep and lasting impact on me. I must also thank Björn and his wife Emma for being two solid and dependable rocks at which I know I can always moor.

Christer Garbis, formerly at the Department of Communication Studies, at Linköping University, sailed a parallel course to mine for a long time. Throughout the years we have met at various foreign ports and I am thankful for certain crucial comments given at some key moments. Among other things Christer has been good at reminding me to occasionally see the bigger picture.

Although we have yet to meet in person I would also like to take the opportunity to thank Klaus Bærentsen for his inspirational work and for a short, but helpful, correspondence.

Thanks also go to Hans Alnemark for helping me along during one of the tougher phases of this project and for doing his job well. Getting a copy of this book into his hands has been part of my motivation towards the end of this project.

The staff at the department of philosophy at Lund University have always been extremely obliging and helpful with practical matters. My heartfelt thanks to Eva Börjesson, Lena Gunnarsson, Jonas Josefsson, Johannes Persson and Ann Tobin for countless things.

I'd would also like to thank the ten people who let me into their homes and allowed me to observe them whilst they cooked. Thanks for patiently answering my strange questions about the minutiae of your kitchens and cooking practices and thank you for ten wonderful dinners. You all contributed to making research the pleasant enterprise that it should be.

A number of foundations and organisations have contributed with financial support along the way. I would like to express my sincere thanks to: Dagny och Eilert Ekvalls premie- och stipendiefond; Stiftelsen Fil dr Uno Otterstedts fond; Stiftelsen fond för Blekingsk hembygdsforskning; and Stiftelsen Syskonen Willers donationsfond. I am especially grateful to The Swedish Foundation for Strategic Research for putting food on the table during a major part of this project.

Special thanks are due Stiftelsen för internationalisering av högre utbildning och forskning (STINT) who financed my sojourn at the Distributed Cognition and Human–Computer Interaction Laboratory at the University of California, San Diego. Thanks also to Stiftelsen Fil dr Uno Otterstedts fond for additional means during that visit.

Ed Hutchins was my host during my six month stay at his lab in San Diego and did much to make my stay enjoyable and illuminating. My sincere thanks to him and Jim Hollan, and everyone else at the lab for their generosity and for making my visit such a pleasant one.

I am also grateful to David Kirsh for inviting me into his lab and for dangling a number of enticing research morsels in front of me. I only wish my time at UCSD could have been longer. Thanks too for some particularly useful comments on paper two of this thesis, and for sharing some valuable thoughts about my approach in general.

Outside of academia my family and friends have all contributed to making life rich and wonderful in a multitude of ways. I am particularly grateful to my wife Charlotte de Léon, who has endured me at my worst and savoured me at my best. Her uncommon ability to talk about the hard things in life contributes to making life that much easier. Thanks for things both big and small. In addition I must thank her for shouldering what became a rather heavy burden towards the end of writing this thesis.

My daughter Miranda de Léon never ceases to bring me joy, and although perhaps not contributing to the timely completion of the present work, I have her to thank for making me occasionally think about someone other than myself.

My in-laws Inger and Per Mellander have helped this thesis along over the years in various practical ways, doing their part in making the lives of me and my family run more smoothly. I must thank Inger, in particular, for the much needed babysitting that she did during the summer.

My pseudo-intellectual comrades in arms, Robert Janulczyk and Fredrik Björklund, have provided occasional catharsis by discussing research and the PhD student's lot with me, which has eased my journey at times. Our discussions about women and sex were generally more fulfilling, however.

I would also like to thank past and present members of my band for allowing me the pleasure of sharing in such a mindless activity, a much needed break from more mindful things; one only wonders why they must also insist on intelligent conversation? Thanks in particular to Stig Berthelsen, Mats Lundblad, and Jonas Mellvinger.

*

And last, but not least, I would like to thank my parents Robert Nordin and Ann-Mari de Léon, not only for that misguided moment of lust in which I was conceived, but for supporting me in everything I have ever done. In the present context it seems especially appropriate to thank my mother for tricking me into reading and for teaching me the basics of good writing, and my father for fuelling my interest in how things work and for doing his best to answer my questions about science as I was growing up. I thank both for making reading, thinking, discussing and writing seem such worthwhile activities, and for discouraging me from getting a proper job. This thesis is, accordingly, dedicated to them.

Lund, August 2003

On Thoughts and Things

Introduction

For almost any activity that you might care to name, there is a great variety of ways in which we might perform that activity. Among the choices open to us are the tools that we use, the ways in which we use those tools, as well as the overall organisation of the activity. But all ways of doing things are not equally efficient, safe, or even fun. Different ways of doing things also tax our mental resources to different degrees.

The tools that we use play a part in this. Any artefact we employ will constrain how an activity is carried out, and thus contribute to structuring the activity in which it is used. Though there is usually some leeway in how an artefact is used, the combined constraints imposed by the artefact and the logic of the task at hand will greatly limit the options.

In the first four of the five papers collected here, I explore some of the different ways in which artefacts can help structure tasks to make them less cognitively taxing, and how these artefacts (and the strategies for using them) develop over time in cognitively beneficial ways.

One of the things that these papers try to demonstrate is how our apparent intelligence often depends on well designed artefacts. It is commonly artefacts "that make us smart" (Norman, 1993). Much, or most, of our intelligence is, therefore, artefactual.

The fifth paper – a co-authored paper – addresses the problem of self-control and how modern sensor and computing technology might be used to support self regulation.

The first half of this introduction presents two themes and situates the present work against contemporary theory. The second half of the introduction gives summaries of the papers collected here.

Three Views of an Artefact in Use

Norman (1991) has proposed a simple, but useful, framework for looking at artefacts in use. I will introduce it here as it will help situate my own work. Norman's framework consists of two perspectives, or views, from which to look at a person using an artefact.

From the "system view" we see the artefact and the person, together, producing some result. The system view is a detached view from outside of the activity. From this vantage point an artefact may appear to enhance a person's abilities. For example, a checklist of the kind commonly used by aeroplane pilots will be seen to enhance memory and performance.

The other view, the "personal view," is one that we can share with the individual that is using the artefact. This is an over-the-shoulder perspective from which we can see how the artefact affects the task to be performed. Take the example of the checklist again. From the perspective of the pilot using this artefact, the checklist replaces the need to remember a sequence, with three new tasks: constructing the list, remembering to consult the list, and reading and interpreting the list (Norman, 1991).

Although the distinction between these two views was originally applied to what Norman calls "cognitive artefacts"(Norman, 1991, 1993; Hutchins 1999) – representational artefacts that "maintain, display or operate on information" (Norman, 1991, p. 17) – I think it can be extended to other classes of artefacts without problem.

To Norman's two views I would like to add a third perspective, what might be called the "artefact view." This view, I suggest, focuses on the physical properties of the artefact and how these properties contribute to the structure and nature of the task. If we were to look at the checklist from this view we would note the physical properties that permit it to play the roles that it does. We might notice how the checklist was written on a two dimensional, flat surface, capable of receiving and retaining an impression (these are some of the properties of transcriptions identified by Latour, 1986). We could then show how the properties of the checklist, seen from the artefact view, relate to the structure of the task, as seen from the personal view. For instance, we could point to how the flatness of the surface of the checklist contributes both to authoring and correct sequential reading of hierarchical lists.

These three views taken together give us an abstract way of looking at people working with things, and we can also start to consider the various ways in which the two perspectives can be related to one another (as was done above in the rudimentary sketch of the connections between the personal and artefact views of the checklist).

Using a particular artefact, or set of artefacts, for a task constrains and influences how a person performs that task. Though the influence of the artefact on the cognitive demands of the task may not always be positive (we have all sworn at unfamiliar ticket vending machines), some tools have a cognitively benign influence on task structure. What interests me in this thesis are artefacts that help make a task cognitively easier to perform.

Cognitive Congeniality

A task may be cognitively demanding for a variety of reasons. It may, for example, be ordered in such a way (or include certain kinds of actions) that attentional and memory resources are overtaxed. Or it might be demanding because of the nature of the cognitive operations involved. It has been noted that we are better at some kinds of mental operations than others. We seem to be particularly good, for example, at recognising patterns, modelling simple dynamics of the world, and manipulating physical objects (Rumelhart, Smolensky, McClelland & Hinton, 1986), and by comparison notably poor at solving abstract logical problems. We are, as Clark likes to say (e.g. Clark, 1997, p. 60), ”good at Frisbee, bad at logic.” Some artefacts help circumvent these kinds of problems by the way in which they contribute to the structure and organisation of a task.

Kirsh (1996) calls the measure of how cognitively hospitable an environment is its "cognitive congeniality." According to Kirsh (1996, p. 440–441) cognitively congenial environments reduce "the number and cost of mental operations needed for task success," "cognitive load on working memory" and increase "the speed, accuracy or robustness of performance." I will employ the term for artefacts that act in a like manner.

The First Theme of the Thesis

Some artefacts contribute to making a task more cognitively congenial. An interesting questions is in what ways they are able to do so? These related questions make up the first of the two main themes of this thesis. The most concentrated treatment of these questions appears in paper two, but these issues also run through papers three and four (see the end of this introduction for detailed summaries of the papers). There are many different ways in which an artefact can contribute to the cognitive congeniality of a task and I have limited my inquiry in a number of ways.

Some Delimitations

In this thesis I have restricted myself to analysing *single* individuals performing *routine* tasks. That I chose to focus on single individuals, acting more or less in isolation, is partly an analytical convenience. This does leave out intriguing issues of communication and co-operation between people acting in groups (and how artefacts might figure in those kinds of processes[1]), but allowed me to instead home in on the fine-grained details of artefacts, to focus on their physical properties and workings and how these relate to the task facing the user (i.e. the task as seen from the personal view). Work focusing on the artefact view of task performance is rare (see summary of paper two for some important forerunners that inspired the present effort) despite an increasing awareness of the importance of artefacts for cognition. Furthermore, as people do work alone at times, limiting the scope in this way did not seem artificial.

1. For some work on the roles of artefacts in the co-ordination of groups working together see e.g. Hutchins (1995a), Hazlehurst (1994) and Garbis (2002).

The other restriction imposed was to focus on routine tasks. Familiar and oft-repeated activities are admittedly just a slice of human activity, and a selection of all of our possible interactions with artefacts. People also encounter and use artefacts daily in both unfamiliar situations and for novel purposes. What little work there is on the properties of artefacts that make them cognitively congenial is mostly concerned with assisting people in figuring out *what* an unfamiliar artefact does and *how* it is to be used. For example, much of Norman's pioneering and highly influential work on the design of artefacts (e.g. Norman, 1986, 1988, 1993) has this emphasis (as he himself points out; see Norman, 1999).

In a routine task, the task is already familiar and the main problem is not to figure out what the task is, or how it is to be executed, but to get through the task in a satisfactory manner (Hutchins, 1990). Routine tasks make up a large part of human activity, but there seems to be much less work in cognitive science and cognitive engineering (for an introduction see Woods & Roth, 1988) on how artefacts can support routine performance. This was one of the reasons for the choice of emphasis in the thesis.

There are several aspects of what makes artefacts cognitively congenial that I have left unexplored. For instance, work on designing the perceived affordances of artefacts (Norman, 1988, 1999; Gaver, 1991), feedback, mapping, conceptual models (e.g. Norman, 1988), or enriching the working environment with clues that aid in appropriate action selection (Kirsh, 1996), are just acknowledged in passing.

Another reason for focusing on routine tasks is connected to the second of the two main themes of the thesis, to which I now turn.

The Second Theme of the Thesis

The defining elements of a routine task are that it is regularly repeated and that it is performed in a customary way. When something is done frequently there is both time and incentive to tweak and adjust the procedure used and the artefacts employed. Almost anywhere we look, where some task or activity is regularly performed, gradual changes occur that, usually, lead to an improvement in the way the activity is carried out.[2] Improvement can be in the efficiency of the task, as in the time taken for the task, the materials used, the energy consumed, the physical or mental effort expended, the likelihood and severity of errors, but also in such areas as how pleasurable the task is to perform and the aesthetics of performance.

In the present context it is those changes in procedure and artefact design which bring along improvements in cognitive congeniality that will be explored. This, then, is the second theme of the thesis and is treated mainly in papers one, three and four.

Though changes to both artefacts and procedures can improve the cognitive congeniality of a task, I will be mostly concerned with the contribution made by artefacts. My realisation that procedures are as important as the design of artefacts for cognitive congeniality (in fact an integral element) has been slow in coming. There are brief glimpses of this realisation throughout (as in papers two and three), but procedures and their relationship to artefact design are not given a comprehensive treatment.

Evolution of Artefacts, Procedures and Tasks

Although some of the changes in artefacts and procedures are due to deliberate interventions, made to improve a task, there are also other, non-intentional, sources of change. Some changes are accidental, or result from external activities that impinge on the task. Some changes even result from the very fact that the task is so often repeated (for example, changes that result from the wear and tear of repetition, or the spatial redistribution of artefacts; for more on this see paper three). Of course only some of these changes will be beneficial (the same goes for intentional design and redesign) and in these cases there is the possibility that a change will be conserved and incorporated into the task.

In this thesis I look at some changes on three different time scales. In paper one we will see changes that can occur within a

---

2. My use of the terms “task” and “activity” is somewhat informal. Usually I mean something vague and broad by “activity” and something more determinate, bounded, and with a specific objective, when I speak of a task. For a summary and comparison of different uses of the concept of a task see Preece et al. (1994, chapter 20).

single hour of cooking. In paper two I have tried to reconstruct the changes in an artefact (and accompanying practices) that occurred over a period of some thirty years. Paper four looks at the changes to a particular class of artefacts (firearms) from the middle ages to the present day, a period of some six hundred years.

For the sake of clarity let me point out that I have not tried to provide a systematic model for how artefacts, procedures and tasks evolve, nor have I tried to link the different time scales to one another (as surely they are linked). The papers simply provide some examples and speculations about these processes at three different time scales.

Artefactual Intelligence

At this point the reader should be ready for an explanation of the title of this thesis. The subtitle should be clear by now, but what exactly is "artefactual intelligence"? The term is used here to capture two different but related notions.[3] The first meaning of the term refers to the intelligence that arises from the *use* of well designed artefacts. This is a property of artefacts and people seen from the system view. The aeroplane pilot, together with his checklist, possesses artefactual intelligence.[4] In this particular example the intelligence of the system is its robust memory and consummate performance.

The second meaning of the term refers to the intelligence and wisdom that can accrue in artefacts over time and which contributes to the artefactual intelligence of the system taken as a whole. This meaning of the term refers to a property of the artefacts themselves. Take the example of the checklist again: the actions and the order of the actions on the checklist reflect a great deal of intelligence and experience that has been "crystallised" (Hutchins, 1995a, p. 95 f.) and made permanent in the artefact.

Artefacts are not the only things that can support, aid, shape or take part in cognitive processes. The natural physical world can

---

3. The term artefactual intelligence was originally introduced in de Léon (1999). The definition given there was, however, rather vague.

4. This use of the term "artefactual intelligence" has the same meaning as the term "cognition" in Hutchins work on distributed cognition (e.g. 1995a, 1995b).

sometimes play these roles, as can social structures and organisations (for introductions to socially distributed cognition see Resnick, 1991; Cole, 1991). Artefactual intelligence is thus a subclass of what some have begun to call situated[5] or "extended" cognition (Clark & Chalmers, 1998).

Situated, Embodied, Extended and Distributed Cognition

The situated view of cognition might be contrasted with a more traditional view of what cognition is and where it takes place. The predominant view in psychology, and related mind sciences, has long been that thought and environment are quite distinct. Once upon a time this was believed to be an ontological separation of mind and matter, now it remains as a tacit conceptual separation, that places all thinking inside skulls and which sees the physical environment as a mostly passive thing, to be acted on and manipulated.

Since about the mid nineteen eighties (but with earlier antecedents[6]) an alternative view of the relationship between thinking and the external world has increasingly gained favour. This alternative perspective recognises the importance of the physical and social environment for cognition. The surrounding world is no longer seen as something simply to be interpreted and acted on, but as something that can support, or even take part in, the processes of cognition.

Though many agree on the basic premise – in the sort of vague formulation given above – there is no consensus theory (as evidenced by the disjunctive heading for this section). Rather than trying to give the history of the field and comparisons and summa-

---

5. I am using the term "situated cognition" as a catch-all term. It actually derives from the very specific position of "situated action" (Suchman, 1987), but it has become almost standard practice to use it like this.
6. Husserl, Hiedegger, Merleau-Ponty, Lewin, Brunswick, Gibson, Barker and Bateson are all people that crop up in the literature, and who have concerned themselves with our involvement with the physical environment. Dewey, Mead and Vygotsky are usually cited in connection with the socially distributed nature of cognition (for additional forerunners see Hutchins, 2001; Cole & Engeström, 1993). Most often the parallels pointed to seem to have been discovered post hoc, and there is rarely a *direc*t connection between these earlier ideas and more contemporary work.

ries of key studies, I will point to what I believe to be some of the key ideas.

Minimal Representational Resources Needed. Partly as the result of work on autonomous robots (Brooks, 1991a, 1991b) and animate vision (e.g. Ballard, 1991), it has become increasingly clear that cognitive agents do not need complete and faithful inner models of the external environment in order to produce adaptive behaviour. It is enough that an agent is able to locate significant cues in the environment and retrieve information from the environment as needed. For example, a person searching for Kodak film in a drug store doesn't have to build a detailed model of the store, but can limit his or her search to the colour yellow. The characteristic colour of Kodak film and advertising serves here as cheap, easy-to-detect environmental cue (the example is from Clark, 1999).[7] Or, once I have determined that you and I are alone in a room I can pick you out as "the large moving object" (McClamrock, 1995, p. 96). In other words, the information needed by the agent is constrained with respect to the amount and kind of information that is required to identify the situation and the objects in it.

Should this strategy fail, the agent can then retrieve more information from the world. Repeated consultations, as opposed to full modelling of the environment, has the advantage of being computationally cheap and fast. The world is after all, as Brooks puts it (1991b), "its own best model."

Evidence for the representational sparseness of human cognition comes, for example, from recent experiments on change blindness. It has been shown that under certain conditions, people fail to detect gross and remarkable changes to scenes and pictures (see the various papers in the special issue of *Visual Cognition;* e.g. Simons, 2000).

Action-Oriented Inner States. Not only does it appear as if cognitive agents exhibit a certain representational frugality – in terms of

7. Also see Swain and Ballard (1992) on identifying objects by their rough colour histograms.

what and how much of the world is represented – but there also seems to be an internal representational economy in which representations (or more non-committedly, inner states) are closely geared to the kinds of actions that an organism "wants" to perform.

This coupling of adaptive behaviour to key properties in the world is nicely illustrated by the cockroach's escape behaviour. The cockroach has two antenna-like structures at the rear of its abdomen that are sensitive to wind fronts. Wind velocities accelerating at 0.6 m/s^2 or more activate escape motions in the roach, lesser wind velocities do not (Clark, 1997). Here is a property of the world that correlates very well with a lunging attacker.

A similar example is the action of the tick. On detection of butyric acid (which is found on mammalian skin) a tick will loosen its hold of the branch it is clinging to and fall on the animal below (von Uexküll, 1934/1992). Here is a state of affairs in the world that permits butyric acid to stand for "mammal" and where the detection of the relevant property immediately, and almost mechanistically, triggers the appropriate behaviour.

Epistemic Actions. Another important notion is the idea that actions, in addition to furthering the physical goals of an organism, can also be performed in the service of cognition. Epistemic actions (as they have been dubbed by Kirsh and Maglio, 1994), can change the state of the world in ways that simplify cognitive tasks. Kirsh and Maglio (1994) list three kinds of epistemic actions.

First are epistemic actions that reduce the memory involved in computation. For example, in a pilot experiment in which subjects were asked to calculate the sum (in dollars and cents) of an arrangement of mixed coins (Kirsh, 1995a), a number of memory saving strategies were observed. One subject, for instance, would "borrow" a nickel from those he had already counted to convert odd valued amounts to even values (thus limiting the phonological complexity of the string to be kept in working memory). He would also keep his thumb over the "borrowed" nickel to save him from having to remember whether a nickel had been borrowed and which nickel to avoid counting.[8]

Then there are epistemic actions that reduce the number of steps involved in mental computation. Imagine the following problem (suggested in Kirsh, 1996). You are faced with a pile of different length sticks scattered on a table in front of you, and you have the task of finding the longest stick. One possibility strategy is to compare the sticks pairwise, picking up two sticks from the table, comparing them and discarding the shortest. By continuing to pick up sticks from the table and comparing them with the one in your hand you will eventually be left with the longest stick.

An alternative way of solving this problem, that cuts down the number of mental (and physical) operations considerably, is to pick up all of the sticks at once and to push their bottoms against the table: the stick that pokes out the furthest at the top is the longest stick.

Finally there are epistemic actions that reduce the probability of error of mental computation. An example of this is a strategy observed in players of the computer game Tetris (Kirsh & Maglio, 1994). The game involves manoeuvring shapes (zoids) that fall from the top of the screen into specific arrangements at the bottom of the screen. Many players will shift the zoids to the far right or left of the screen, and then back again, to count out their column placement and reduce errors in spatial judgements. This action momentarily takes the person away from the physical goal, but has a crucial cognitive function.

Extended Cognitive Processes. Another daring notion is the idea of cognitive processes extending into the world, and parts of the world having functional roles usually ascribed to the brain (such as storing and transforming information). A nice example of information looping out into the world is the control mechanism for the flight of the common fly. Apparently (Marr, 1982), when the fly takes off, there isn't a signal sent from its central nervous sys-

8. For other examples of information being encoded in the environment see e.g. Kirsh (1995b, 1996) and Beach (1988). These recount strategies that mainly make use of spatial location. Another pervasive strategy is, of course, to make use of inscriptions.

tem to its wings. Instead, there is a direct link between the feet of the fly and the flapping of its wings, so that when its feet lose contact with a surface the wings start to flap. To fly, the fly first has to jump, and a signal is then sent from its feet to its wings!

One of the clearest and most thoroughly worked out version of this basic idea is Hutchins' theory of distributed cognition (Hutchins, 1995a, 1995b). Rather than speaking of individual cognition looping out into the world, Hutchins takes a step back and analyses the cognition of systems comprised of people working together with artefacts. Hutchins does not deny individuals cognition, but shows that certain systems, such as the navigation bridge of a large ship (Hutchins, 1995a), can be given a cognitive analysis. The navigation of the ship can be seen to be accomplished by the flow and transformation of representational states through both people and artefacts.

In the last paper of this thesis my co-authors and I argue that the same kind of distributed analysis can be given to the "hot" motivational and emotional processes as have been given such "cold" cognitive phenomena as reasoning (Zhang, 1997), memory (Hutchins, 1995b) and collaboration (Hutchins, 1995a; Rogers & Ellis, 1994). In the paper we also discuss some of the ways in which people distribute motivation with the help of other people and artefacts.

The above brief sketches capture what I consider to be the core insights and themes of situated cognition.[9] Needless to say, these are ideas and hypotheses that are not shared by everyone working in the field, and among those who subscribe to these ideas (myself included) there is usually little explicit mention of these concepts. A reason for this, I think, is that so much work in the area takes the form of ethnographically inspired field studies, with video

9. For some very readable overviews of the current state of situated, distributed, or extended, cognition see Clark (1997, 1998, 1999). Kirlik's introduction to applied work on cognition and the designed environment (Kirlik, 1998) is also very good, and unusual in that it includes mention of human factors work, HCI and cognitive engineering. It is hard to delimit what to include in "situated cognition," but activity theory should also be mentioned as it seems to share many of the same concerns and aims (for some parallels and comparisons between distributed cognition and activity theory see Nardi, 1996; Kaptelinin, 1996).

recordings and dialogue transcripts making up the bulk of the data. The analysis of these studies usually occurs at a different level of description, for instance in terms of dialogue turns, cultural models, actions, strategies and the propagation of representational states. In the future I fully expect to see more work that crosses these boundaries of scale, level and style.

The Historical Basis of Cognition

According to theories of situated cognition we owe part of our cognitive achievement to our physical and social surroundings. Social and artefactual structures usually have a long history which underlies their present shape and form. For example, most of the artefacts that we use have gone through a series of designs and redesigns and a number of earlier uses. There is a very important sense in which the capacities of a person using an artefact derive, in part, from changes that have aggregated over time.[10] If we subscribe to some version of situated cognition, then asking questions about the genesis, growth, development and appropriation of these kinds of structures should be as natural as asking questions about their current roles.

Though many of the authors cited in the section above *do* acknowledge that there are interesting processes responsible for the build-up of cognitively significant physical and social structures, these processes then figure to a negligible extent in their accounts.

Hutchins (1995a), for instance, gives truncated histories of the astrolabe and the compass rose (both ancient navigational instruments that significantly transformed the cognitive task of ship navigation), but then simply concludes that practice can be "crystallised" into things, without discussing the process of crystallisation itself.

Kirsh (1995) acknowledges that the interaction of agent and environment can be studied along different time scales, and does an admirable job of looking at medium and short term mecha-

10. The idea that our environment is suffused with the achievements of prior generations in material form has its roots in the works of Hegel, Marx, Dewey, Durkheim, Leont'ev, Luria and Vygotsky (Cole & Wertsch, n.d./1999).

nisms of how people set up their workplaces for particular tasks, but the issue of how the workplaces evolve is not addressed.

Bærentsen's (1989) work on the evolution of the rifle is an inspirational exception which explicitly deals with the interaction of artefacts and the cognitive demands of task performance, as well as the influence of cognition on artefact development. Although a bold and innovative attempt, Bærentsen's analysis relies on a problematic notion of cognitive processes being "built into" things (for a critique see the fourth paper of the thesis). Bærentsen's paper is written in the activity theory tradition, which places great emphasis on the historical and cultural foundation of thought and artefacts. This should, therefore, be an ideal place to find the kind of analysis sought for. The activity theoretical concepts of *externalisation* and *historicity* also seem to capture the concerns discussed. However, as Engeström (1999) has noted, there seems to be a general paucity of work in activity theory on these very topics.

Summaries and reflections on the papers

In the rest of this introduction I summarise each of the papers briefly and try to highlight what I believe to be the main points and contributions. The papers were written over a span of several years and some are closer to the present time of writing than others. Where sufficient time has passed I am sometimes able to offer my more recent thoughts on a particular paper and the issues treated by it.

Paper One – Actions, Artefacts and Cognition

The first paper of the thesis details a cognitive ethnography (cf. Hutchins, 1995a, 1995b; Lave, Murtaugh & de la Rocha, 1984), which I conducted in the year 2000, of people cooking in their homes. It is the last paper actually written, but is based on data gathered near the beginning of my project. My observations, as well as a preliminary analysis, lay dormant as I worked on and wrote the other pieces of the thesis, but informed most of my subsequent work (in particular papers two and three).

The choice of cooking as a domain of study was motivated by the general dearth of work on productive tool-use, and on isolated individuals. Most of the ethnographic field work in situated and distributed cognition has been conducted on people working in groups, and usually working with external representations (such as maps, diagrams, texts, computer displays and other kinds of inscriptions). Hutchins' work on ship navigation (Hutchins, 1995a) and the piloting of aircraft (Hutchins, 1995b; Hutchins & Klausen, 1996) clearly falls here, as does work on air traffic control (e.g. Halverson, 1995) and rescue management (e.g. Garbis, 2002).

An advantage of studying groups (and group-work involving the production and manipulation of external representations) is that the researcher has better access to the processes of the group, compared to the relatively shuttered mind of a discrete individual. The communication between members engaging in joint work opens up a window through which the researcher can gaze. When the task performed by the group is a socially distributed cognitive task there is a very real sense in which it is possible to "step inside the cognitive system" (Hutchins, 1995a p. 129).

It is interesting to note, therefore, that one of the rare studies of individual tool use, Keller and Keller's study of artisan blacksmiths (Keller & Keller, 1993, 1996), makes use of a method involving participation and self report (see Keller, 2001). This is not the only method that they use, and people working singly can be studied without this kind of direct involvement.[11]

Several of the observations reported in Kirsh (1995b), taken from an unpublished study of people cooking, show people using the spatial arrangement of objects to simplify choice, perception and internal computation.

The examples from cooking that were given in Kirsh's (1995b) paper, together with my interest in the cognitive aspects of tool use, prompted me to conduct my own study of cooking. One of the objectives of my study was to collect further examples of actions, strategies and artefact use that contribute to making tasks

11. Note, however, that any study of situated practice requires profound domain knowledge. Analyses of situated cognition require interpretation of actions and speech acts of the people under study. Often, a method of systematically excluding alternative interpretations is involved in analysis. Theoretical domain knowledge, and personal experience with the activity analysed, are prerequisites.

cognitively congenial, and perhaps also to validate some of Kirsh's findings.

At this time I had also begun to think about the historical dimension of cognition (and had written a paper that touched on this theme; see paper four below), and so I was hoping to catch a glimpse of the processes by which tools are adapted by people (in ways that improve the cognitive congeniality of task performance). The domain of cooking was promising, as most people have a long history of cooking experience, and the task of cooking a meal is sufficiently complex for there to be strong incentives to simplify the activity.

My study was explorative and not directly tailored to collecting data on the build-up of strategies and physical structures. Despite this I managed to collect some very nice data on the accumulation and appropriation of cognitively significant physical structures. There is, for example, one observation in the paper that shows a woman figuring out how to extend the use of a prefabricated measuring cup, to handle a situation that occurred whilst baking a cake. In the paper I argue that having established this new use for the measuring cup, the woman in the study lay the ground for future cognitive savings.

There are also other observations that show how cognitively beneficial arrangements can be generated as a by-product of the repeated performance of certain activities. For example, the kitchen implements that are most frequently used are also most likely to have been recently washed. There is therefore a good possibility that commonly used items will be on the plate rack drying. The processes of using and cleaning kitchen implements sorts out, and makes readily available the most frequently used implements, making them easy to find and use. Some items may never be put away, but lead an existence either in use or drying on the rack. This was the case, for instance, with one of the participants' kitchen knives.

One of the subjects I visited had a considerable collection of spices and was able to recount several phases in the development of this collection. As this data was very rich I conducted some follow-up interviews. This part of the study can be found in paper three.

Another section of paper one recounts a number of observations that are familiar from Kirsh's (1995b) paper; behaviours that utilise the spatial layout and arrangements of objects to encode information. In this part of the paper there are examples given of people preparing and maintaining their work spaces.

Another group of observations in the paper concern some ways in which participants handled timing constraints. There were also some suggestive, and perhaps surprising, observations of clock use, in which clocks were studiously consulted at the beginning of a task, but then failed to figure in determining the end point of the task.

Cooking is a rich domain. It takes place in a highly structured environment and involves a number of tools and implements. There are complexities of timing and co-ordination, but at the same time there is a great deal of flexibility in how constraints are handled. We may be tempted to think of cooking as a simple task, but that is an illusion born of our prolonged experience with cooking – it is easy to forget the multitude of meals that a normal adult has prepared.

Neither is cooking restricted to the kitchen. The preparation of a meal may begin much earlier, as when we prowl our local delicatessen and discover what's available today. Choices made in the shop reflect and echo choices made in the home.

Furthermore, since cooking is such a frequently recurring activity there is time and incentive for people to hit upon, learn or invent new tricks and techniques, and to change the organisation of their kitchens. The kitchen promises to be a good place to find special solutions tailored to problems and tasks that are frequently encountered. In the present study I visited each participant and filmed them cooking only once. It would be interesting to revisit the same group of people to observe them cooking on several occasions, to observe them cooking familiar recipes as well as novel ones and to see them cooking various kinds of meals.

Paper Two – Cognitive Task Transformations

This paper outlines a number of principal ways in which artefacts can reduce the cognitive burdens of performing routine tasks. The focus is on single actors working on familiar tasks, where the main difficulty is one of satisfactory performance rather than that of figuring out *how* a task is to be executed. Starting from a commonly held assumption that tools transform the structure of the tasks in which they are employed (e.g. Cole & Griffin, 1980; Hutchins, 1990, 1995a; Norman, 1991), I explore some basic ways in which tools can assist in reshaping tasks to circumvent excessive demands on attention, working memory and motor control, and the need to perform the kinds of cognitive operations at which we are biologically less suited.

The principal task transformations that are identified include: removing superfluous actions from a task; delegating problematic parts of a task to other people or to artefacts; substituting less demanding parts of a task for the parts that are cognitively demanding; rearranging a task so as to avoid conflict between certain actions; and increasing the tolerance of a task environment to sub-optimal task performance.

How these transformations were arrived at is not mentioned in the paper and it might be appropriate to do so here. Each transformation was distilled from examples taken from my own experience, the field study of cooking that I conducted (see paper one), and other published accounts of people interacting with artefacts. For each case that I considered I tried to determine how the artefact in question changed the task at hand and also how it achieved those changes. I suspect that the transformations singled out in the paper are not exhaustive and that there may be examples which could be given that overlap the categories proposed.

Although it is widely recognised that artefacts can restructure tasks, I believe the present systematic effort to be unprecedented.

A limitation of the approach taken in the paper is that cognition is reduced to a fairly simplistic resource model in which memory and attention are taken to be more or less fixed quantities, irrespective of the particulars of the task being analysed. Although a more veracious conception of cognition may yield slightly different results (and a revisitation of these issues may very well be warranted), it is an interim simplification that has permitted me to make a number of significant distinctions and to discern some interesting patterns.

The transformations described in the paper have sometimes been mistaken for a proposal for a formalism or as a framework for task analysis. Yet another concern voiced is that breaking down tasks into linear sequences of discrete actions, as is done in the paper, is unrealistic and overly deterministic.

To clarify: the intent of the paper was assuredly not to provide a formalism (and although pretty, the schematic figures in the paper should not be mistaken for one), nor to propose yet another task analytical framework (for examples of these see Diaper, 1989; Hackos & Redish, 1989; Stammers, Carey, & Astley, 1990). It is rather what might be called a "broad brush" approach to interaction analysis (Green & Petre, 1996).

Neither do I think that real life tasks are rigid and unchanging. The fact that they do change is, after all, one of the main themes of this thesis. I will insist, however, that actions usually follow other actions, like "beads on a string" (Gatewood, 1985). So even though there is likely to be flexibility, variation and creativity even in routine tasks, we can still profitably look at short sequences of actions and analyse the fine details of how artefacts shape and structure those short sequences. It is this attention to the fine details of the interaction with an artefact which, in my opinion, is one of the main points, and strengths, of the paper.

Considering the physical properties of artefacts, and how these influence (constrain, make possible, support or guide) the activities in which they figure, was partly inspired by work by Latour (1986) on the properties of transcriptions in relation to the making of science, and Gedenryd (1998) on the properties of the materials used by designers that make them suited for various cognitive aspects of designing, and certain passages in Hutchins (1990, 1995a), in which the roles of a couple of artefacts in the task of navigation are analysed, as well as the pioneering work of Norman on cognitive engineering (e.g. 1986, 1988).

In paper two the focus is limited to the ways in which the physical design of artefacts can transform tasks, but in my study of cooking (papers one and three) it is clear that actions, procedures and techniques can often do much the same kind of thing. That techniques can alter task structure is mentioned in passing in paper two, but could have been elaborated. Clearly, for several of the categories proposed in the paper there are plausible analogical transformations that do not require any change in the physical props but only in technique or procedure.

Above I pointed to some of the limitations of the approach taken in paper two, as well as some possible sources of error. Furthermore, I tried to allay some potential misunderstandings. I'll end this review and reassessment of paper two by pointing to what I believe to be some of its merits.

First of all, the paper provides a framework that may have some utility when comparing different tasks, especially when those tasks are alternative means of accomplishing the same, or similar, end. The taxonomy and the figures provide us with a means of abstracting from the complexity of task performance, as well as a succinct way of visualising the processes of people interacting with things. In addition to furnishing us with new analytical tools, the paper also serves to clarify some of the specific roles that artefacts can play in activity and outlines how it is that they can perform those roles. As such, the paper is a step towards an understanding of the extended mind and the basis for artefactual intelligence.

Although I am wary of claiming relevance of the framework for the design of artefacts and task environments – the application of theory to design having a notorious track record (see e.g. Caroll, 1991; Landauer, 1991) – it could potentially play *some* such role. Perhaps the framework presented here (together with other work on cognitive support, e.g. Walenstein, 2002) could help "raise the level of discourse" of designers (Green & Petre, 1996). The framework might be employed to highlight certain difficulties encountered in a task, aid us in understanding those difficulties, and helps us in the search for alternatives.

Paper Three – The Cognitive Biographies of Things

In paper three the idea that we might construct "cognitive biographies" for things is proposed and developed. These cognitive biographies are accounts that detail the history of an artefact and its use and focus on the physical changes undergone by the artefact over time, as well as the cognitive corollaries of those changes. Calling these use-histories "biographies" captures something of their intimate relation to particular individuals, and also gives an indication of their temporal scale.

There are a number of reasons for why we might want to construct such biographies (why it might even be incumbent on us to do so). First of all, if we subscribe to an extended view of cognition in which the material realm is deeply embroiled in much, or all, of cognition, then the genesis, evolution and adjustment to cognitively significant physical structures must be seen as an integral aspect of our cognitive achievements. The histories of things and their use must be an essential part of any complete account of cognition.

Even without a commitment to an extended or situated view of cognition, a full appreciation of the history of a task and the tools for that task may also be required if we are to discern the cognitive roles *currently* being played by an artefact; this may be of some interest in, for instance, the process of designing or redesigning artefacts. Against the backdrop of earlier incarnations of an activity and previous forms of an artefact, the cognitive functions of a thing are more easily discerned. By overlaying succeeding phases of an activity with preceding ones, areas of possible cognitive significance can be highlighted and explored.

The idea that cognition has a historical basis has a long and fine pedigree. It is therefore somewhat surprising that these kinds of processes – beyond rather broad and large-scale generalisations – have received so little attention. Part of the reason for this is, no doubt, the difficulties inherent in observing, or reconstructing, what are primarily historical processes. Unless we limit our interest in the ways in which artefacts and practices co-evolve to very short time scales (for a nice study in this vein see Agre & Shrager, 1990), we have to choose between longitudinal studies and historical reconstructions, both beset by some worrisome methodological concerns.

Of the two options, a longitudinal approach may, initially, seem preferable: the reliability of the data gathered is less open to question and we seem to avoid much of the speculation and conjecture required in a reconstruction. But longitudinal studies demand great effort without guaranteeing results, and we cannot beforehand know that anything of interest will turn up at a chosen site or during the time-frame selected. In many cases the sheer scope, and also intrusiveness, of longitudinal projects make them unfeasible. However, there may very well be a productive middle ground to be explored. Some of the suggestions for future work that were given above, in the section on paper one, would surely fall here.

Although less controlled and more speculative than a longitudinal study might be, reconstructions permit us to explore sites where the occurrence of significant artefactual change has already been established. They also present an opportunity to investigate real-life events and changes that span long time periods: durations for which a longitudinal approach would be impractical. These particular characteristics make reconstructions (or cognitive biographies, as they are called in the paper) a compelling and intriguing possibility. It is an exploration of a reconstructive approach that I opt for in paper three.

In this paper I take a shot at constructing a cognitive biography of an artefact and its use over a period of roughly 30 years, exploring the mutual influences between cognition, activity and changing physical structures. The artefact in question is an unusually large spice shelf that I encountered whilst conducting the cognitive ethnography detailed in paper one.

In the resulting biography I start from a description of the contemporary setting and current use of the shelf. This is based on video data of the shelf being used in a single session of cooking. Video and an interview then formed the basis for a cognitive analysis of the ways in which the physical structure of the shelf, the jars stored there, and their organisation, presently supports some of the cognitive work of cooking.

An account is then given of the history of the shelf leading up to its present shape and use. The history begins with a single shelf containing just a few conventional Western spices, takes a course via a small box used in a period between more permanent residences, and ends up at the large, highly structured, shelf that is used today.

Whereas the history of the shelf is based on interview data, the various ways in which it has been used, in its different phases, is mostly supposition. Again, parallel with the history, is an analysis of the probable cognitive consequences of each change to the shelf.

A number of supplementary interviews were conducted after the initial study in order to flesh out the history of the shelf and to check some of my interpretations.

In the process of reconstructing this cognitive biography some important insights were made along the way. One is the realisation that much of the structure that supports cognitive activity may actually have partly *non-cognitive* origins. It seems as if at least some structures are the result of chance, circumstance, compromise, surrounding agents, as well as the shaping force of repetition.

Not only is this an important corrective to theories of design that are overly intentional (for a critique of some of these see Gedenryd, 1998), but it also says something unexpected about the source of our cognitive powers. We humans have a remarkable capacity for creating artefacts – tools and structures that both supplement and extend our cognitive powers. It is a capacity that has often been attributed to an uncommon ability to envision and foresee, but is perhaps equally derived from a talent for conserving and adapting to fortuitous structural changes.

Another important, and related point, is the significance of use. One thing that is demonstrated in the paper is how the cognitive congeniality of an environment is as much a function of an agent's particular use of that environment as it is a function of the environment itself. It is the particular ways in which things are used that *permit* them to contribute in cognitively beneficial ways. This serves to underline that cognitive congeniality is a relational property, and that it must be studied as such.

The paper is a kind of experiment. By having written it I hope to have shown what a cognitive biography might look like and what kinds of things we might learn by constructing them.

The method employed in the paper has its evident weaknesses, the most obvious being the amount of speculation required in constructing a biography of the kind given. However, as I point out towards the end of the paper, these problems are ones that are shared by a number of enterprises involved in reconstructing changes in material culture (its uses and its meanings). I suggest that it might be profitable to pool resources with these other fields (e.g. archaeology, the history of technology, the anthropology of technology, social construction of technology studies and actor-network theory) and to share methodological innovations and insights. More immediate dividends, however, may be had if we begin by exploring a middle ground, somewhere between constructive cognitive biographies and lengthy longitudinal observational studies.

The kind of study I have in mind consists of regular ethnographic work, over a restricted period of time, of a person (or persons) working in some artefactual environment. A suitable site would be one in which we have reasons to expect changes in the working environment, as well as changes in strategies and procedures. Remaining within the domain of cooking, we could, for instance, study the progression of someone becoming familiar with a particular recipe, or we could follow the establishment and commencement of a new restaurant. A suggestion that came up in conversation with Jean Lave, is to study someone moving house. If we limit ourselves to cooking, for the sake of the example, we could study the present use of the kitchen, the process of packing, dismantling and moving to the new home, and the process of setting up and adjusting to the new circumstances. Needless to say, we are not limited to the domain of cooking; another kind of situation, that immediately occurs to me, is that of a person setting up in a new office.

These examples of possible future work have all been given rather elliptical descriptions. They could, and should of course, be more fully worked out. In the present context they are intended to wet our appetites.

Paper Four – Building Thought Into Things

This paper is the fourth of the thesis, but was actually written before the others. Thematically it seemed appropriate to place it here, however. In the version that appears here I have allowed myself the luxury of correcting a few grammatical eyesores and a few instances of unclear reference that can be found in the published version of the text. I will quickly recapitulate the papers that precede it to show how it fits into the thesis as a whole.

Paper one (the study of action and tool-use in cooking) described some of the complexities and subtleties of people interacting with things, and paper two (on cognitive task transformations) proposed some ways in which to analyse and understand some of those interactions. The next paper (paper three, on cognitive biographies) also focused on the cognitive aspects of interacting with things, but paid particular attention to the ways in which an artefact and its use changed over time, and to the cognitive consequences of those changes. The time scale of paper three was roughly that of a generation and the focus mostly limited to a single actor and the practices and artefacts tied to that one actor.

Artefacts also change over longer time scales, sometimes over several life spans, and for many things there is a long history and ancestry: a line of previous forms and associated practices. In the fourth paper of this thesis the evolution of a particular class of artefacts, over a period of several hundred years, is explored, as well as the cognitive consequences of these changes. One might say that in moving from paper three to paper four the thesis moves from the cognitive biographies of things to their genealogies.

Acting as backbone for paper four is a piece by Bærentsen (1989) in which the evolution of the rifle – from the middle ages to the present day – is examined from a cognitive perspective. Bærentsen's central thesis is that the actions and thought processes required to operate a rifle at a particular stage of its historical development are "built into" subsequent generations of the artefact. This claim itself is not without precedent (it is, for instance, one of the tenets of activity theory, the tradition in which the essay is written), but the consistency with which Bærentsen pursues the thesis, as well as his attention to detail, is inspiring.

Much of my own paper is spent reassessing Bærentsen's analysis and, in the process, gradually arriving at a rather different conception of the relationship between cognition and things. Along the way an attempt is made to find a non-metaphorical way of making sense of the notion of thought being "built into" things. I am, however, unable to find a viable interpretation of the concept.[12]

12. That I named my paper "Building thought into things" was unfortunate as this has misled a number of people into believing that I am advancing the very thesis that I am actually arguing against.

The alternative view that emerges is one in which particular attention is paid to the interplay between the changing physical properties of artefacts and the structure of the tasks in which they are used. To a certain extent this is simply a rediscovery and corroboration of the, by now, familiar view that artefacts shape the structure of tasks (e.g. Cole & Griffin, 1980; Hutchins, 1990, 1995a; Norman, 1991; and the second paper of the present thesis). However, the analysis provided in paper four elaborates on this basic idea in a number of ways.

One of the virtues of the paper is that it tries to go beyond simply acknowledging the connection between things and performance, to actually explore some of the details of that relation. This is my early attempt at connecting the physical properties of artefacts to how they make the transformation of a task possible – a line of inquiry that eventually became paper two. For example, the ability of some artefacts to "retain state" (such as the steady combustion of a burning fuse, or the force stored in a spring) is linked to the timing of a task, and some mechanical features of medieval firearms are linked to the way in which attention is employed when using these artefacts.

In the paper, the transition from one stage of an artefact to the next (in this case, the transition from one kind of rifle design to the following generation of rifle design) is described as a number of interrelated transformations. That I chose to call them all "transformations" is regrettable as it implies a commonality between them that does not in fact exist. The transformations actually touch on different aspects of the transition, at different levels of analysis, and on different time scales.[13]

Ill-chosen nomenclature aside, the paper details some key aspects of what happens in the move to a new artefact design. Invariably, changes in the design of an artefact are followed by changes in its use as well as in the ends to which it is put. Sometimes the development in the design of the artefact leads to tasks that are cognitively less demanding for the user. The elegance of a

13. Incidentally, these transformations are only remotely related to the transformations described in paper two. This is, potentially, another source of confusion for the reader.

well designed artefact is that it presents the user with a simple sequence of actions to perform, whilst ensuring that that sequence still meets the requirements of the task. To do so the artefact has to help bridge this gap by taking the simple actions of the user and "translating" these into acceptable output.

There are at least two ways in which an artefact can help bridge the gap between the demands of the user and the demands of the task. Take, for example, a door handle. A standard door handle permits certain actions (such as those movements that conform with the direction in which the handle is hinged) and resists others (such as attempts to pull the handle outwards or upwards). Movements that are oblique to the rotation of the handle, however, will be shaped, or guided, to conform with the affordance of the object. There is a sense then in which the physical properties of an artefact may, in addition to affording certain kinds of action (Gibson, 1979; Norman 1988), also partly shape or transform those actions.

Not only may an artefact tolerate imperfect handling (for more on this see paper two), but in some cases it also transforms the user's actions into task relevant output. Take the example of the door handle again. The reason why pushing down on the handle opens the door is that the mechanism of the lock, as a whole, transforms the movements of the user in a way that retracts the bolt from the door jamb.

This is a direction of influence, from artefact to person, which is rarely taken into consideration. Also, thinking about the artefact as something that has to solve two simultaneous problems – that of meeting the requirements of the user and the demands of the task – may be a useful way of thinking about artefact design. We are used to thinking in these sorts of terms when dealing with the design of user interfaces for computers, but the present discussion shows that we may think about other types of artefacts, including the mechanical ones analysed in paper four, in much the same way.

The paper's major fault is, perhaps, that so much time is spent on Bærentsen's thesis and on retracing each step of his analysis. As Bærentsen's paper is only available in Danish this does provide a wider audience with the story, but I could have focused more on my own contributions. Writing the paper was an exploration conducted more for my own benefit than for the edification of potential readers. Working my way through Bærentsen's analysis was an important step for me, that helped me formulate some of the themes that run through the rest of this thesis, as well as some of the solutions proffered.

Another thing that is missing in the paper is a discussion of how cognition is involved in the evolution of artefact design. The paper deals with the cognitive consequences of artefact change, but it would also have been interesting to have analysed the contribution of cognition to those changes. We can speculate, for instance, that when cognition that has been tied up in some part of a task is suddenly freed, it can be put to new use. As an artefact and task evolve perhaps "cognitive slack" (to coin a term) is quickly taken in and put to use elsewhere. This is an issue that would have been worthwhile to explore.

Although there are surely more pleasant artefacts than weapons to analyse, the evolutionary history of firearms and the practices surrounding them are ripe material worthy a revisitation. My more recent, but brief, forays into military history and weapons technology (e.g. Dupuy, 1984) have revealed a wealth of further details, nuances and complications that just beg to be analysed. I would like to think that the analytical tools provided in paper two and three would be particularly useful in such an enterprise.

Paper Five – The Future of Self-Control

The topic of the last paper of the thesis (a paper co-authored with Lars Hall and Petter Johansson) is the perennial problem of self-control. In our paper we present a basic model of the domain of self-control and then provide a range of suggestions for how modern sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities. The model consists of two core concepts. The first we have called computer-mediated extrospection (CME). This concept builds on the familiar idea of self-observation or self-monitoring, and concerns itself with the crucial need for accumulation and explication of self-knowledge in any rational person-centred decision process. The second concept is that of distributed motivation, which was mentioned earlier in this introduction in connection with distributed cognition. This concept also ties in rather nicely with, and extends, the idea of pre-commitment and self-binding, that is often discussed in the self-control literature.

In its simplest form the problem of self-control consists of the fact that we tend to choose smaller, sooner rewards, rather than larger, later rewards, despite knowing that this is against our best interest. In a calm moment of reflection we may decide on a certain course of long term action, but then temptation crosses our path and we are easily swayed. For example, most of the people who decide to quit smoking make that decision based on the expected long term health benefits. The irrationality of succumbing to the lure of smoking, after having made such a decision, is evidenced by the fact that the person is capable of acknowledging their long term goals before, during and after the illicit cigarette has been consumed. The feelings of regret that often follow on such lapses further underscores the self-defeating nature of failures of self-control.

For most people, relying solely on internal self-regulation is seldom enough. But, in the same way that we can make use of the environment to alleviate the cognitive demands of everyday life, so might we use stable social and artefactual structures to compensate for "weakness of will." Indeed, if we look at what people actually do, at least some efforts seem to be devoted to processes that distribute the burdens of sustained motivation out into the world.

There seems to be an abundance of such culturally evolved and personally discovered strategies, skills, props and tools that people use. People write authoritative lists and schedules, put salient markers and tracks in the environment to remind them of appropriate actions, rely on push and pull from social companions and family members, latch onto role models, lock themselves into institutional arrangements, seek out formal support groups, and even hire personal pep coaches.

We believe that these kinds of strategies can be analysed as distributed phenomena in a similar way to how distributed cognition has handled cognitive processes. We therefore invite ethnographically inspired field work on distributed motivation and look forward to an expansion and adaptation of the distributed cognition framework.

Distributing motivation into the world is a partial solution that certainly helps, but which does not ensure success; failures in self-control are too numerous to mention (see the paper for alarming statistics on some select self-control issues). The general problem with these distributed strategies is that they do not have enough binding power. Take the example of a person wanting to exercise on a regular basis. That person's family, for instance, may be willing and able to encourage that person to exercise, assist by making time available, and even censure the family member when he or she fails to follow through on the decision. The problem is when the person in our example stops going for their daily run, or stops going to the gym. It is all to easy to make excuses and exceptions, lie or simply shrug off reminders and jibes.

Another example that illustrates the problem is an often told anecdote about a young Afro-American man who made a commitment to pay USD 20 to the Ku Klux Klan every time he smoked a cigarette. In contrast to the case above it is easy to understand the force this commitment *might* have had on his behaviour. The fact still remains, however, that once he has succumbed to the temptation of smoking a cigarette, nothing really compels him to transfer money to the KKK. In fact, it would be irrational to pay the money. But if no such crucial deterrent for future behaviour can be established, then why on earth should he adjust his behaviour in relation to the commitment to begin with? If he could exert the type of mental control that would effectively bind him to pay the smoking fee to the KKK, then why not simply bind himself not to smoke in the first place?

One of the solutions that we propose in the paper is to offer people the opportunity of increasing the binding power of the precommitments that they make. We suggest a number of related and intertwined strategies that employ modern sensor and computing technology for their effect. As an extreme example, imagine the force that the young man's precommitment would have had if the forfeit was the automatic transfer of funds to the KKK from the young man's bank account. The severity of such a scheme would probably be too much, so much that people would refrain from entering into it in the first place.

The schemes that we offer try to achieve a gentler balance. The model proposed (see figure 1 in paper five, p. 150) starts with techniques for handeling easy problems and progresses toward more difficult problems of self-regulation. The first schemes all revolve around presenting a person with information that is not otherwise readily available.

Towards the more gentle end of the scale we propose systems that simply help people maintain important goals in an active state. In the face of distracting (and often tempting) stimuli it is easy to loose track of the goals that we have set ourselves.

Further along the scale are context-aware systems that can recognise user behaviours indicative of impeding breakdowns, or that react when they find the user in a specific "context of temptation." These systems can provide a person with early warning.

Next are systems that help people to better appreciate how they are progressing relative their goals. A straightforward application would be a context-aware device that counts the amount of calories or cigarettes consumed. However, these systems promise even greater impact in relation to goals that are more abstract and distantly long-term. For example, imagine someone who has decided to become a more amiable and caring person. By operationalising this goal, and monitoring goal progression, that person could be provided with discriminating feedback on the outcome of her behaviour.

In the rest of our model we present some different kinds of precommitment technologies that allow a person to bind themselves to varying degrees. At one end are schemes that utilise short self-binding, or micro precommitments, an example being a remote control that you can use to turn off other appliances for a set limit of time. Micro precommitment technologies can help us overcome transient temptations. At the extreme end are schemes that offer almost irrevocable binding. Take the following, real, example. Some larger casinos give patrons, who are prone to too much gambling, the option of having themselves banned from playing. Since casinos are generally equipped with rigorous security and surveillance systems, the ban can be very effectively enforced.

This is a long and rich paper and I have not been able to do it justice in this short summary. The paper identifies an important class of problems suitable for ameliorative action using sensor technology and ubiquitous computing (a field currently preoccupied with solving technical issues and suffering from an absence of significant applications), it proposes a number of new theoretical constructs as well as a model, and even suggests new avenues for ethnographic research.

Lastly, let me try to dispel a possible misconception. The proposal of using sensor and computing technologies to alleviate problems of self-control may frighten some readers. A scheme such as this can all too easily conjure up images of surveillance and control. Our paper ends with a discussion of the issues of flexibility, and ethical concerns about privacy and persuasion, but I would like to stress a central aspect of our perspective.

The future technologies that we envision are intended to allow people to enter into precommitments on a *voluntary* basis. There are also different strengths and lengths of precommitment and we foresee systems that can be tailored to, and by, the users themselves.[14] Although the opportunities and possibilities for coercion will increase with the rise of networked technologies (a seeming inevitability), the scheme we propose is one aimed at empowering the user.

14. In this respect we feel that we differ somewhat from the neighbouring field of persuasive computing (Fogg, 1999, 2000, 2003). Though we certainly share many of the goals of persuasive computing (such as the concern for environmental, health and safety issues), and even some methods, we reject its persuasive element and pre-packaged morals.

PAPER ONE

ACTIONS, ARTEFACTS AND COGNITION

AN ETHNOGRAPHY OF COOKING

ABSTRACT: This paper details a number of observations from an ethnographic study of ten people preparing a meal in their own kitchens. The focus of the study was on how people cope with the cognitive demands of a familiar task such as cooking.

The observations of the study are grouped under three headings. One section describes a few ways in which some of the participants handled timing constraints, as well as some perhaps surprising observations of clock use. The next section recounts some ways in which the spatial layout of objects was used to encode information (cf. Kirsh, 1995a), and also some examples of preparation and maintenance of the work space, that arguably benefit cognition. The final section concerns how the working environment and tools are adapted and adopted. One segment of data shows how the use of an already present artefact can be extended, other data points to cognitively beneficial structures that are generated as by-products of the repeated performance of cooking.

INTRODUCTION

In the last ten years or so, a number of etnographically inspired studies have ably demonstrated the deep contextual nature of cognition. Studies of grocery shopping (Lave, 1988), ship navigation (Hutchins, 1995a), the piloting of aircraft (Hutchins, 1995b; Hutchins & Klausen, 1996), commercial trawling (Hazlehurst, 1994), the use of office technology (Suchman, 1987), and dairy workers assembling product orders (Scribner, 1986), all show various ways in which cognition can lean on the physical and social surroundings.

It has been shown, for instance, that people can make use of the world to remember things (Hutchins, 1995b; Beach, 1988; Norman, 1988), to simplify choice, perception and internal computation (Kirsh, 1995b; Clark, 1997), and to transform tasks to make them less cognitively taxing (Hutchins, 1990, 1995a; Norman, 1991; de Léon, 2002).

Cognition as it is portrayed in this body of work is heavily reliant on, or even partially constituted by, external artefacts and social interactions. Whether the external world is viewed more as scaffolding for individual cognition (see e.g. Salomon, 1993), or as a near equal partner in a distributed cognitive process (e.g. Hutchins, 1995a), all would agree that the material and social context must figure in a satisfying explanation of human activity.

The fieldwork cited in the first paragraph (with the exception of Lave, 1988) has been conducted in settings were groups of people work together. One of the advantages of studying people in groups is that the communication between individual members are observable. Since much of human work and activity occurs in groups this choice of focus is also natural.

Another common characteristic of the domains above is that they involve work of a mathematical or computational nature. Lave's grocery shoppers (1988) compute best buys, Hutchins' navigators collaborate in computing their ship's position, Scribner's dairy-workers assemble and total the sum of orders. These kinds of tasks are ideal for study as they are usually more tractable than such tasks as the production of goods (say, weaving or pottery).

Many of the tasks commonly studied also involve external representations, like maps, diagrams, texts and computer displays. Work centred around external representations is a central part of much of human activity and, again, provide the methodological advantage of more observable phenomena.

For these combined reasons little work has been conducted on single individuals engaged in productive tool-use. Keller and Keller's work on artisan blacksmiths (Keller & Keller, 1993, 1996) is an exception that deals with both single actors and tool use. The main cognitive factors that their work focuses on are the roles of imagery, and how knowledge is used in the preparation for, as well as the engagement in, productive tool-use.

Kirsh's paper (1995b), on the intelligent use of space, also focuses on single users and on the manipulation of objects. The paper contains several interesting ways in which the spatial arrangements of objects can be used to simplify choice, perception and internal computation.

The Study

Inspired by the examples given in Kirsh's paper, I conducted an ethnographic study of ten people cooking. One objective was to collect further examples of actions, strategies and artefact use, that contribute to making a task more "cognitively congenial" (Kirsh, 1996), and perhaps to validate previous findings (such as those in Kirsh, 1995b). I was also interested in observing how people cope and co-ordinate with the various demands placed on them whilst cooking.

A related aim was to try to catch something of the processes by which tools are adapted by people in ways that ease the burdens of task performance (cognitive and otherwise). Cooking takes place in a highly structured environment and involves a number of tools and implements. Since it is a recurring activity and many of the constituent parts of cooking are often repeated, the kitchen promised to be a good place to find special solutions tailored to problems and tasks that are frequently encountered.

METHOD

The participants recruited for the study were all people who volunteered to take part and were not compensated for doing so. During recruitment the participants were informed that I wanted to study them in order to better understand how they worked when they cooked in their kitchens.

Ten people took part: five men and five women. With the exception of one participant (a man in his mid fifties) all were around thirty years of age (see table 1 for the exact ages of the participants and for a summary of other personal characteristics). The sessions all took place in the participants' own kitchens. They all cooked alone, with the exception of two sessions in which the participants' children where present.

Name	Gender	Age	Occupation	Meal	Familiar w. preperations	Recipe used
Amanda	F	31	Book editor	Dinner for 4	Yes	Yes
Annabel	F	34	Librarian	Dinner for 6	Yes	Yes
Belinda	F	32	Lawyer	Dinner for 4	No	Yes
Benny	M	28	PhD Student	Dinner for 4	No	Yes
Elisabeth	F	25	Secretary	Dinner for 3	Yes	No
Henry	M	30	PhD student	Lunch for 2	Yes	Yes
John	M	35	School teacher	Lunch for 4	Yes	No
Lisa	F	33	PhD student	Lunch for 2	Yes	Yes, for dessert
Marcus	M	29	School teacher	Lunch for 3	Yes	Yes
Robert	M	55	Social worker	Lunch for 3	Yes	Yes

Table 1: A summary of the pseudonyms used, other personal characteristics of the participants, and the conditions for each session.

The participants were asked ahead of time to cook something that they would ordinarily cook on the day on which I visited them (though some clearly made slightly more elaborate efforts than usual). In half the sessions the meal prepared was a lunch time meal, the other sessions were all dinners.

Each session was video filmed with a small hand held video camera at medium range and lasted for about an hour; the duration of the sessions was dictated by the time it took for each subject to prepare the meal. The participants were asked to think-aloud as they cooked and I would also occasionally ask for clarifications. All but one of the participants spoke Swedish (one spoke Danish) and all transcripts have been translated into English (by the author). Back translations, for the purpose of validating the translations, have not been conducted.

After the sessions the participants were interviewed about their cooking histories. All were prompted to talk about how they began to cook and about when and how the implements in their kitchens were acquired. The interviews were not filmed. Instead written notes were taken. The format of the interviews was informal and different kinds of ground were sometimes covered with different participants.

In the post-cooking session I asked to be given a quick tour through the cupboards and drawers of the kitchen. The tour was also video taped. On a couple of occasions brief pencil sketches of the layout of the kitchen were made to help me later when reviewing the tapes.

In one case, which contained an extremely rich material on the build-up of supportive structures, a number of supplementary interviews were made after the initial session, and additional digital still pictures taken (see de Léon, 2003a).

The rest of the paper presents a collection of observations taken from the data. These have been loosely grouped into three main themes. The first one concerns some ways in which time can be handled in a session of cooking, specifically how people cope with misalignments in timing. This is followed by a collection of strategies that all utilise the visible grouping of objects (or the removal of visible distractors) for their effect. Included here are strategies that make use of spatial placement (and orientation) of objects to encode task relevant information. The third and final section of the paper deals with some different ways in which the physical structures in a kitchen change over time, as well as how people can expand the uses of artefacts that do not change.

TIME AND TIMING

A number of observed behaviours can be collected under the rubric of time and timing. Time is important in almost any activity and cooking is no exception – this is no doubt confirmed by the reader's own experiences in the kitchen. Not only do various parts of a task *take* time, but different parts often have to be co-ordinated with each other (for instance, the timing of one part may be dependent on that of other parts). In addition, there are sometimes deadlines set by factors that are extraneous to the activity itself, as in the case of a meal that has to be ready by the time guests arrive or a particular TV show begins.

One way to meet these various demands on time and timing is to plan the activity in some detail before it takes place. Though people surely do plan at least some of their activities (or parts thereof), other kinds of strategies also seem to be at work. Next I will recount an episode which shows how an activity can be set up so as to allow room for improvisation, thus obviating the need for detailed planing. This section will be followed by one in which some ways of coping with misalignments in timing are outlined.

Preparing to Improvise

One of the participants, Belinda, had chosen to cook a stew that she hadn't prepared before, but had eaten at friends. She began her session by cutting up pork and some spicy sausages. She reasoned as follows:

Belinda [Cutting up some thin sausages into small segments]
I'll wait with the rest of the ingredients and fry these up first. I think I can just put the rest in the stew, not much cutting to do

Moments later she added the following:

Belinda *I'm doing this*
[Moves to stand by the stove]
before, because I think the rest can cook whilst the rice is going

What Belinda did was to reason that most of the necessary preparations could be done whilst the rice cooked. There were some parts of the meal, however, that she felt were more time consuming, and so she did these first (i.e. cutting up some sausages and pork and frying them). With these preparations out of the way she felt confident that she would be able to improvise through the rest of the session (or, at least, that she would have time to handle any eventualities that might arise).

Benny too made a decision much like this one, comparing the time that would be required to prepare some rice with that required to prepare a chicken:

Benny *It doesn't take very long to prepare the rice so we'll do the chicken first*

Clearly Belinda and Benny are both engaging in a form of planning, in the broadest sense of the word: both assess certain future actions to be taken and decide on an order in which to perform them. However, the "plans" created are very general and specify broad aspects of the activity, and not the fine details. By identifying more time consuming processes of their activities, and engaging in these early on, Belinda and Benny increase the likelihood that improvisation will be a viable strategy.

Strategies That Compensate for Imperfect Timing

Just because a person chooses to improvise doesn't mean that the constraints inherent in an activity disappear. Even if an activity has been planned down to the smallest details there will still be situations in which a particular task is completed either too early or too late.

When cooking at home there is usually no need to be maximally efficient (though cooking for a large dinner party may be a different matter), imperfect temporal alignment of various activities are tolerated and adjusted to. Nevertheless, there are tasks for which timing does matter or end results will suffer. In John's session, potatoes that had been set to boil were ready too soon, well before the rest of the meal. Allowing the potatoes to continue boiling would result in potatoes that end up too soft. Taking them out of the pot early would allow them too cool. Instead John utilises a simple compensatory strategy:

John [Looks about, then reaches for an oven potholder. Lifts the lid off the pot containing the potatoes and pokes the potatoes with a fork]
The potatoes are soon done here. I think. I think they are
[Puts the lid back on the pot]
What you then do is to turn off that hot plate
[turns off the hot plate]
It's not quite done yet the potatoes, but almost. It can stand there until they're ready

The potatoes are almost ready, but are not in phase with the rest of the meal. By taking them off early and leaving them in the hot water the process of cooking is slowed down to allow the rest of the meal to "catch up." John need not pay them any further attention until the end of the session. Taking them off a little early reduces the risk that they will be over-cooked. Belinda does something similar partway into her session:

Belinda [Is standing by the stove, then moves to the work bench]
And I'll do some foil over the meat
[Opens a drawer and takes out some aluminium foil]
should have done that before

	[Pulls off a sheet, returns the packet of foil to the drawer, then moves back to the stove and covers the meat dish with the foil]
Author	*Because*
Belinda	*So they don't loose too much heat and they can lie here and cook a bit more kind of slowly kind of*

This general strategy (or strategies) of preparing something in advance, or preparing something partway, is so pervasive that it is often factored in at the beginning (i.e. planned for).[1]

In the examples above a preparation is completed too early and has to wait for the rest of the activity to "catch up," as it were. Another possible strategy is to try to catch up with event that are about to pass one by. An example follows taken from Robert's session:

Robert	[Sprays some oil into the bottom of the pan. Picks up the small ceramic bowl with spices and tips them into the pan. Places the bowl back on the cutting board. Picks up a wooden spoon lying next to the stove, drops it, and picks up one lying next to it. He stirs the spices into the oil and sprays more oil into the pan. He returns to stand by the cutting board and moves the empty bowl out of the way. (The spices can be heard crackling loudly in the simmering oil.) Working quickly he splits the onion in two with a kitchen knife and puts the knife to the side. He picks up a cleaver and continues to chop the onion-half, first one way, then the other, and finally the bit remaining at the tips of his fingers. The bits of the chopped onion are scraped to the top edge of the board] *They can't be allowed to burn that is why I'm hurrying a little*

1. The microwave oven might be mentioned here, not as a fast and convenient way of heating food, but as a device that allows one to compensate for misalignments in timing when cooking.

[Robert cuts the second half of the onion. First one way, then the other, and then chops the last piece]
I'm getting a little stressed, I see it is smoking from here
[Robert takes the cutting board and holds it over the pan, scraping the bits of onion off the board and into the pan. He stirs onion and oil with a wooden spoon]

In this example the problem was that one activity (the chopping of an onion) was not completed in time (before the oil began to simmer). Robert's solution was to speed up that activity. Some other possibilities might have been to take the oil off the hot stove, until the onion was ready, or to settle for the quantity of onion that had already been chopped at that point. These are other possible solutions and we can't know for sure why he chose as he did. We can note, however, that the action actually chosen was one that allowed the activity to retain a certain flow and momentum.

A combination of the two main strategies outlined above seem to be at work in an episode in Marcus' session. Marcus too is boiling potatoes and decides, as the potatoes are almost done cooking, that he needs the hot plate they are standing on for something else. Apparently one of the stove's other hot plates works less well and Marcus wanted to move a pot of peas to the efficiently functioning plate the potatoes were on:

Marcus [Lifts the pot with potatoes off the stove and places them on the work bench. He puts the pot of peas on the hot plate the potatoes were standing on]
Shall we try them just in case
[Takes the lid off the potatoes and prods four of the potatoes with the cake tester]
Yes they are soft, they can stand there and keep warm

Moving the potatoes to the side and making room for the pot of peas is a means of catching up: of reducing the delay in the preparation of the peas. But the potatoes are ready and Marcus lets them stand in the hot water of the pot to keep warm, just as John did. Since the potatoes are already done, a downside to this strategy is that the potatoes may end up slightly over-cooked. John, who utilised a similar strategy, took the potatoes off the stove a little before they were completely done.

Clock Time

Above were some strategies that compensate for imperfect timing. Let me underline that strict timing is only important some of the time: for much of cooking there is a great deal of flexibility as to how and when things are done, and the goals of cooking may even, on occasion, be renegotiated.

But more precise timing can be a desirable thing, which is part of the reason why clocks and egg timers are such prevalent artefacts in kitchens. It is therefore interesting to note that in the ten sessions recorded neither timers, clocks nor watches were extensively used. Six of the participants made use of them (Belinda, Elisabeth, Lisa, Marcus, Henry and Benny), three did not, and one case is uncertain due to a break in the taping of the session. In those cases where a timer, clock or watch (and in one case, a mobile telephone) was in fact used it was sometimes employed otherwise than expected.

For a task like cooking rice or boiling potatoes we might expect a watch or clock to be consulted and an end time calculated. We might, furthermore, expect the person to periodically consult his or her watch until the calculated end time drew near. However, two of the sessions (Belinda's and Marcus') contain episodes in which clocks were studiously consulted to begin with, but nevertheless failed to figure in determining the end points of the tasks engaged in.

I will begin by giving you a number of consecutive segments from Belinda's session. In this case (and the next) I have included the starting time for each segment:

00:00 [Standing by the stove she pours a carton of cream into the pot]
Whilst that's simmering we'll start the rice as well

02:20 [Stirs the pot]
We'll start the rice
[Bends down, pulls out a drawer, takes out a pot and a lid]

04:50 [Pours water from the tap into a one litre measuring cup]
Six decilitres of water but I put in some extra
[Holds the cup near eye level and looks at the scale. Walks over to the stove. Tips the water into the pot were the rice has been frying for a while in oil]

08:35 [Walks over and grabs the kitchen roll. Glances towards the rice as she walks back to the work bench. Starts pulling paper off the roll and holds her arm up to look at her wrist watch]
I want to have some idea of the rice. It's a quarter to
[Starts to wipe down the cutting board with the kitchen paper held in her hand]
about another quarter of an hour

13:58 [Takes the lid off the rice, puts it back on]
It's coming along just right

16:35 [She is standing by the stove. A possible glance over at the rice]

20:30 [Standing by the stove, glances over at the rice. Picks up an oven potholder. Lifts up the lid and looks inside. Puts the lid back on]

23:30 [Starts walking towards the stove]
So now I have to check the rice which I think is about to

[Hesitates, turns around, looks down at a drawer, opens it and takes out a fork. Walks over to the stove. Picks up the oven potholder and takes the lid off the pot in which the rice is cooking. Accidentally drops the fork, picks it up again. Tastes some of the rice, puts the lid back on the pot. Takes the lid off again]
Mmmm
[Tilts the pot, looks inside, puts the pot down. Places the lid on the stove. Picks up a one litre measuring cup. Starts walking to the sink]
Too much
[Fills the measuring cup at the tap. Returns to the stove and pours some water on the rice. Replaces the lid]

30:40 [Screws the cap back on the bottle of olive oil held in her hands. Reaches over to the stove and turns the burner off]
I'm turning off the rice now, I think its done

The above is a fair length of transcript, all to make the following rather simple observation. Belinda consults her watch and works out when the rice ought to be ready, but then proceeds to check the rice three times, at roughly five minute intervals, until she decides that it is ready, all without consulting her watch ever again. So why did she look at her watch and calculate an end time to begin with? I'll return to this question shortly after having first related a similar (and similarly lengthy) sequence of segments off Marcus' tape, an episode in which he is boiling potatoes:

00:00 [Looks at the clock]
Twelve past twelve, usually needs about twenty five minutes

08:12 *Hasn't quite started to boil properly*
[Turns the knob. Checks on the potatoes. Looks over at the wall clock]

It's twenty past. Needs at least another twenty minutes

10:32 *Now they're coming along nicely*
[Turns down the heat]

12:37 [Lifts the lid of the pot of potatoes]
Potatoes are cooking good

19:02 [Takes the lid off the pot and puts it on the stove whilst looking at the potatoes]
Try them a bit
[Opens a drawer, locates a cake tester, brings it out and skewers one of the potatoes: lifts it up then lets it slide off the cake tester. Prods two more in the same manner and puts the cake tester down]
Probably need, six seven minutes something will be right
[Puts the lid back on the pot, but leaving a small gap]

23:08 *I think I'll check on the potatoes again*
[Takes the lid off with his left hand, picks up the cake tester with the right hand and prods four potatoes]
It's actually getting there
[Puts the lid back and puts the cake tester down on the bench]
a minute or two until I turn off the hot plate

24:06 [Touches the pot containing the potatoes and adjusts its position on the hot plate]

24:25 [Takes the lid off the potatoes and looks at the boiling water. Puts the lid back. Turns the heat down]

27:20 *I'm checking the peas*

[Takes the lid off the pot of peas]
It's not very good this hot plate so it takes a long time to bring the water to a boil
[Puts the lid back on]
What I usually do then is to swap with the potatoes since they're done now
[Lifts the potatoes off the stove and places them on the sink. Puts the pot of peas on the hot plate the potatoes were standing on earlier]
Shall we test them just to be sure
[Takes the lid off the potatoes, picks up the cake tester and prods four of the potatoes]
Yes they are soft they can stand there and keep warm

In the first segment Marcus looks at his wall clock and establishes the starting time for his potatoes. He also reasons about how long the potatoes will need to cook. About eight minutes later he checks on the potatoes and discovers that they haven't started to boil, he then looks at the clock again and decides on a new end time. This is the last time that he looks at the clock. Much like Belinda, he then proceeds to check on the progress of his potatoes at fairly regular intervals.

I find these two sequences interesting in their own right, but will indulge the reader and proffer some speculations about what might be going on. One possible interpretation of their behaviours is that they initially set out to use the clock, to help them determine when the rice and potatoes were going to be ready, but then forgot to do so or were side-tracked. Though possible, I find little to support this interpretation (though note that warrants for *any* interpretation are hard to establish). Neither of them mentions their failure to keep track of clock time, which might be expected (and was also commented on by both Henry and Annabel when this was the case for them). Both also embark on close and regular monitoring of the rice and potatoes almost immediately.

Perhaps looking at the clock and reasoning about the end time, as both did, served some other function than marking the time at which rice and potatoes would be ready? Looking at a watch or clock may be a way of taking stock, of reflecting on what one has just been doing, and on what might be done during the interval specified. By thinking about the activity in terms of time, one shifts to a more abstract perspective that might allow other kinds of reasoning than a more episodic conception of the activity would. Thinking about boiling rice or potatoes as something that takes a fixed number of minutes is a succinct way of representing that activity and of reasoning about what might be done in that time slot. In addition, by looking at the clock it is also possible to further abstract the activity and to reason about time spatially. Various tasks can then become spatial segments of the clock face that can be superimposed, combined and rearranged on the face of the clock (see Hutchins, 2002, on conceptual anchors).

Perhaps the instances of clock use recounted above are carried over from other cases of clock use in the same setting. When cooking pasta or eggs precise timing may be more appropriate, and even necessary, than when one is boiling potatoes. Looking at the clock may simply be a habit, something which seems reasonable to do, and which incurs little cost, but which then plays no real part in shaping the activity.

The suggestions above are speculative, but do suggest that ethnographic study of the *actual* use of clocks, watches and timers might be productive. The suggestions given here are all possible hypothesis that could be investigated.

VISIBILITY

A number of observations have been collected under the rubric of “visibility,” as they hinge, in some way or other, on the construction of visible groups of objects (or the removal of visible distractors).

Spatial Coding of Information

Kirsh’s seminal paper from 1995 describes, among other things, some ways in which people use space to encode various aspects of a task (also see Beach, 1988; Scribner, 1986). The placement of objects can, for example, encode a person’s location in a task, as well as what action to take next. Several of the examples given in that paper (Kirsh, 1995b) are taken from a similar study of cooking. Some of the same behaviours can be found in my own data. The next two examples, for instance, show how space can be used to encode category membership. The first of these shows spatial location being used to keep track of which knives have been sharpened and which are still to be done. In the second example both location and orientation of mushrooms are used to distinguish rinsed mushrooms from unrinsed ones:

Robert [In front of Robert is a work bench and a cutting board. On the right hand side lie three kitchen knifes and a potato peeler. A cotton place mat is also lying on top of the knives. To the right side of the board is a whetstone. Robert picks up the place mat and tosses it aside and then moves a can of beer out of the way. He picks up the whetstone with his right hand and transfers it to his left. The right hand then picks up the rightmost knife. He sharpens the knife and then places it to the right of the cutting board. He picks up the next knife, sharpens it and places this one to the right of the knife already put aside. He picks up the last knife, sharpens it and places it to the right of the board and to the left of the previous two knifes. He then picks up the potato peeler and moves as if to sharpen it, then puts it down to the left of the cutting board. The whetstone is put away in a drawer near the work bench.][2]

Marcus [There are eight large mushrooms in a loose pile on the cutting board. He picks up the rightmost mushroom and rinses it under the tap. The mushroom is returned to the board and placed upside down at the same time as the next one is picked up. The second mushroom is rinsed and replaced

2. At the same time as he is sharpening the three knives Robert is also conversing with his wife. The conversation has been omitted for reasons of clarity.

> upside down next to the first, just as the third one is being picked up. This one is rinsed and replaced upside down. The procedure is repeated with a fourth mushroom. As the fourth is being returned to the cutting board the other hand picks up the fifth and sixth mushrooms together and then rinses them. They are replaced upside down simultaneously as the seventh and eight are picked up together in one hand, rinsed and then replaced upside down]

That Marcus works two-handedly contributes to the efficiency with which the task is carried out, but also helps to ensure that rinsed mushrooms are kept separate from unrinsed ones. The two handed action, the fact that Marcus picks mushrooms from the left side of the board and works his way leftwards, and that rinsed mushrooms are replaced on the board upside down all provide cues to their rinsed or unrinsed state. Another reason for placing them upside down is that this also makes it easier when they are later cut into quarters.

Here it is difficult to separate out the cognitive motivation for some action, or actions, from other requirements of the task. It is unlikely, however that actions are (as perhaps implied by the distinction in Kirsh and Maglio, 1994) either "pragmatic" (i.e. done to bring an agent closer to some end state) or "epistemic" (i.e. done for cognitive reasons). Actions may serve both of these functions simultaneously.

In the episodes recounted, spatial grouping obviate the need for the participants to remember which knife has been sharpened, or which mushrooms have been rinsed. This information about the objects is instead given by their spatial location and, in the case of the mushrooms, also their orientation. Another example in which spatial location indicates identity is found in Benny's kitchen:

Benny [Backs away from the sink where he's working]
Salt
[Turns and looks down at the cook book]
maybe /inaudible/ *how much salt there should be*

	[Reads the recipe. Returns to the sink and picks up a small metal tin, removes the lid. Puts the lid back and then replaces the tin. Picks up an identical metal tin next to it]
	I have the same tins
	[takes the lid off the tin]
	That isn't so smart, especially if you have sea salt since that looks like sugar
	[displays the contents of the tin]
Author	*How do you know what's what then*
Benny	*I have the sugar there*
	[points to the table behind him. Laughs. Pours some salt into a mortar]
	It took me quite a while before I thought of it
Author	*Having them in different places*
Benny	*They will get mixed up otherwise too. But otherwise it's simple to see the difference between salt and sugar but sea salt gets so, the same lustre kind of*

Keeping salt and sugar separate would seem to be a relatively simple problem. And of course it is. However, it is not the only problem facing Benny when he cooks, but one of many things he must keep track of. Add to this the potentially quite unpleasant consequence of mistaking salt for sugar and it is clear that this kind of spatial coding serves him well.

Another participant, Henry, kept his spices in a number of small clear plastic bags inside a cupboard. There weren't many of these bags to keep track of, but during the tour of Henry's kitchen I asked him how he could tell the spices apart:

Author	*How do you know which spices are which. Do you recognise them*
Henry	*Often … I will recognise*
	[Lifts a bag to eye level]
	I can see this one is thyme for example
	[Puts the bag back inside the cupboard]

	But then there is a problem if I have to decide between saffron and curry, then you have to taste
Author	*And you do*
Henry	*Yes and if I'm lucky it's saffron*
Author	[laughs]

Clearly, even a few spice jars (or bags) are enough to keep track of. It is therefore perhaps no great surprise that the enormous collection of spice jars found in Robert's kitchen is so extensively structured (see de Léon, 2003a).

Preparing and Maintaining the Work Space

The strategies delineated above work best if there are no distractors present, no superfluous objects competing for attention. Clearing stuff away at the beginning of the session (as Amanda, Annabel and Benny all did), or during breaks in the flow of the activity (as did Annabel, Benny and Lisa) improves the cognitive congeniality of the work space. As Benny told me:

Benny	*I always try to keep it clean. You know, free surfaces in some way so you know what you're doing*

And as Amanda put it:

Amanda	*I can't do this if there are too many things* /out/

As well as removing possible distractors, clearing away clutter may also be performed in an effort to be tidy, to keep the workspace clean, or simply to provide space to work in. Perhaps it is this mixture of factors – cognitive and pragmatic – that motivates Annabel:

Annabel	*I like to wash up kind of bit by bit, I find it difficult when everything is out*
Author	*Difficult how*
Annabel	*Here it's mostly that there is no work surface, I don't know, there is something, it's sort of cluttered and difficult*

Preparing and maintaining the work space by removing distractors and ensuring that all visible objects are relevant to the task at hand is one way of simplifying the cognitive work of cooking. Another strategy employed by some of the participants was to start their sessions by bringing out all the ingredients they would be using. As Marcus explained:

Marcus *Then I usually bring out most of what I'll be using, so I kinda have an idea, this gives me a certain overview of what order I'm going to do things*

If we take Marcus at his word, then it seems as if the physical presence of the ingredients to be used assist him in preparing or planning for the task facing him. And surely, with the ingredients visibly present he doesn't need to remember what they are and they can prompt his memory as he reviews the coming task.

That the mere presence of the ingredients play some such role for Marcus is supported by the following occurrence in which Marcus is confronted with some extra vegetables that were brought out by mistake at the beginning of his session. These vegetables had been purchased at the same time as the other ingredients for the meal, but were not part of the meal being prepared:

Marcus [Looks down at the cook book and reads]
Eh
[Looks up at the two peppers and the leek placed in front of the book, then back at the book again]
These weren't for this recipe, but for another one
[He picks up the two peppers and the leek and goes to the fridge and puts them inside]
that we also bought stuff for when we were shopping

If Marcus was not relying on the purposeful presence of the ingredients about him, then why should he be momentarily thrown by the presence of an additional leek and some peppers? To deem them irrelevant he had to consult the recipe.

Bringing stuff out before it is to be used is a strategy that can also be combined with some of the spatial strategies discussed in a previous section. The following is an example from Amanda's session in which the ingredients for a spaghetti sauce that she is making are grouped on one of her cutting boards before being put to use:

Amanda [Picks up the plastic cutting board resting on top of the wooden one and puts it into the sink. Walks past the sink to the cupboard. Takes out a bottle of ketchup and a bottle of wine. She holds these objects under an arm and opens the refrigerator (which is placed inside a cupboard) and takes out a tube of tomato purée. She closes the cupboard door, walks back to the wooden cutting board and places the three items on it. Walks to a hanging cupboard, opens it, moves a packet of crisp bread aside. Takes out a small packet of stock.]
Meat stock
[Tips out a cube from the packet. Puts the small cardboard packet back in the cupboard and closes the door. Walks with the cube of stock to the wooden cutting board and opens the wrapper against the surface of the board. Opens a drawer under the work surface. Takes out a garlic press and closes the drawer. Picks up a clove of garlic and puts it into the press, then places the press on the board]

Bringing out all ingredients (as well as tools and other items) at the beginning of the session is a way of delegating some aspects of remembering to the world, but it is also simultaneously a way of ensuring, before the task is fully underway, that everything that will be needed will be available. Like Marcus, Benny too starts by bringing out the items that he'll be using:

Benny [Places cook book by window]

	It's a small kitchen so you have to arrange things quite a bit to
	[Goes to fetch bag of groceries standing by the door to the kitchen. Places bag on kitchen table]
	I'm otherwise relatively structured when I cook
	[Takes two packets of risotto rice and places them on the table]
Author	*Is that because the kitchen is small*
Benny	*Yes, no, I also think it's because I think… I usually take everything out that I'll be using*

Benny then continues to remove items from the bag, naming each one as it is put down. And at the end he notices that something is missing:

Benny *I actually forgot the rosemary*

So, preparing the work space in this way – clearing away superfluous items and bringing out everything that will be used – allows objects and implements to serve as reminders of things to be done. Furthermore, they can also be arranged in the kinds of ways described in the previous section. Provided that things are also cleared away during the course of cooking, the absence of ingredients will assist the judgement that the task is complete.

Development of supportive structures and strategies

An assumption underlying the present study was that the physical environment and available materials are often intimately involved in the processes of cognition. A number of ways in which tools and strategies can support the cognitive work of cooking have already been presented, as well as some other features of action and activity, but the question of how materials come to have the shape and use that they do, and how strategies are acquired, has so far been left unaddressed. This is the topic of the present section of the paper.

Sometimes new structures are made or introduced into the kitchen, sometimes old ones are modified, sometimes structures change through use, and sometimes it is a change in the way in which something is employed that permits it to shoulder some of the cognitive burden of cooking. Though the introduction of new structures, or modification of ones already present, is often outside the time frame of a single session of cooking, there were some clear instances that will be related below.

Lets begin with an episode in which an already present structure was co-opted by the participant and transformed into a cognitive resource.

Appropriating Already Present Structures

In the example to be given bellow one of the participants, called Lisa, found a new way of using one of her kitchen implements. Having settled on this new use for the object, and having even established its appropriateness, the artefact in question will simplify certain tasks for her in the future. The example is thus given as a case in which a material resource is appropriated, in a way that allows it to improve the cognitive congeniality of a certain type of task.

When I visited Lisa she had decided to bake a cake to top off the meal she was about to make. At one point in making the cake the recipe specified a certain quantity of coconut flakes. The quantity was given in grams in the recipe, but Lisa's kitchen was without either scales or conversion tables. What she did have was an old glass measuring cup marked down the side with scales for the weight equivalencies of rice, sugar and porridge oats, in addition to markings for metric volume. The cup is a material resource that enables weight to volume conversion to be carried out, for a number of common ingredients, using some rather simple physical manipulations.[3] In the example given below Lisa uses the cup to solve an immediate problem facing her, but in so doing expands her future use of this particular material resource:

Lisa *Coconut flakes*

3. A neat feature of the cup is that volume to weight conversion is carried out simultaneously and instantaneously with the act of measuring out a desired quantity.

[Looks at the bag in her hand. Turns her head towards her son, who is sitting on her hip, and speaks to him in a falsetto voice]
120 grams
[Looks at the work bench]
the scissors
[Lifts a kitchen towel lying on the bench]
where are the scissors
[looks around]
there
[picks the scissors up and cuts off corner of the packet of coconut flakes. Picks up the measuring cup, looks at it whilst turning it. Continues to look at it. Glances at the packet in her hand. Puts the measuring cup down, glancing at it as she does so. She takes some coconut flakes from the bag and gives her son to taste. Picks up the measuring cup. Looks at it and turns it in her hand]
Then we have the next problem. Can't measure coconut flakes
[Puts the measuring cup down on the bench]
So
[Holds up the bag. Looks down at the measuring cup]
We'll say that they weigh, weigh like porridge oats maybe
[Turns the bag in her hand]
Yes that'll do
[Bends down and pours coconut flakes from the bag into the measuring cup]
they weigh like porridge oats
[Moves to put the bag down, but then lifts it up to eye level. Turns it over in her hand]
200 grams. We should be using half the bag
[pours a few more coconut flakes into the cup and then empties the cup into a pan. She then pours out more flakes and empties the cup again]

It is clear, both from what Lisa says, and from her uncertainty in how to proceed, that Lisa has no predetermined strategy for dealing with the problem at hand. There is an artefact available that is clearly made for the job (or, at least, for jobs like it), but it doesn't quite fit her needs. What Lisa does first is to make a judgement regarding the similarity of the density of coconut flakes and the available alternatives. Having done so she then checks the soundness of her decision, of equating the density of oat and coconut-flakes, by using the material resources at hand. The bag is printed with the weight equivalence of almost twice the quantity of coconut flakes specified by the recipe and by seeing that the bag was half empty Lisa was able to conclude that she had achieved an acceptable approximation using the cup.

Of course, had it occurred her, she could have employed this particular strategy – of using half the bag's contents – to begin with, without having to bring the measuring cup into play. But having established the approximate equivalence of porridge oats and coconut flakes, as she did, she not only solved the immediate problem, but also expanded her future use of this particular material resource.

Supportive Structures Generated Through Use

In the above example the artefact in question was already present in the working environment and remained unchanged throughout the encounter. The key transformation, in this case, was in Lisa's understanding and use of the measuring cup. Next, I would like to give some examples in which it is the world that changes (in ways that, arguably, improve the cognitive congeniality of the environment), but where no concomitant conceptual change is supposed or required.

The chief organisational mechanism I wish to bring to the fore is one in which the repeated performance of an activity shapes the surrounding environment in ways that later supports more of the same activity (cf. Barker, 1968). As examples, think of the footpath kept clear through use, or the discolorations caused by tools and implements hanging on the walls of a workshop that facilitate their correct replacement.

In the present study it is not wear and tear, as such, that generates supportive structure, but rather the spatial redistribution of artefacts and other objects as they are handled in various ways. Take the following example from the very beginning of Henry's session:

Henry	[The door to the fridge is open]
Author	*Let me see your fridge*
Henry	*Okay,* [laughs]
Author	*Do you have a special, system*
Henry	*In my fridge? Hmmm, everything that is old and gross is at the back*
Author	[laughs]
Henry	*The further out you go the newer it is*
Author	*Is that so*
Henry	*It is. It's not so long ago that I had a fridge…* [Picks up a plastic bag from the bottom shelf with some cheese inside] *here is an example of, we can throw it away right away* [Henry turns to the sink, opens the door under the sink and throws the cheese in the trash]

The conversation quoted above is light and humorous, and it is quite possible that Henry is partly joking. For one, there is obviously more order to the items stored in his fridge than he himself allows. One of the shelves on the inside of the fridge door, for instance, houses a row of standard sized milk containers and three glass bottles of ketchup, the shelf above it holds some smaller glass bottles, and the topmost shelf inside the fridge contains items that are visibly taller than the other items in the fridge. So contrary to Henry's claim, there is some order, order that is imposed by the size of items and the various spaces that are available. The remaining contents of the fridge, however, does seem to lack any readily apprehensible order. At the end of the session Henry confirms the account he gave above:

Henry	[Opens the fridge door]

	So it's a lot like I said... yes... Hmmmm
	[He reaches into the fridge and touches a plastic bag of potatoes on the middle shelf]
	that one I wouldn't want to swear on how old it is for example
	[Touches a packet of bacon, then a bag of potatoes on the bottom shelf]
	old potatoes and
Author	*So the old stuff is at the back, or*
Henry	[Takes out a small plastic box containing cherry tomatoes, turns it at an angle, looks at it, puts it back]
	since I push it in after awhile
	[Makes a few pushing gesture with the flat of his left hand]

In addition to the order imposed by the sizes and shapes of various food items, as well as the different storage spaces of the fridge, items are also arrayed roughly in accordance with when they were purchased and/or last used. Provided that Henry knows when something was bought, or last used, this emergent order will help him in locating sought for items. The acts of placing, removing and replacing food in the refrigerator thus result in an unintended order that supports the process of locating specific items.

The shelves of Henry's cupboards were also organised according to frequency of use (as were Benny's and Belinda's): items that were most often used were, according to him, placed low in the high hanging cupboards of his kitchen, and thus more easily accessible. Though the data doesn't reveal how his cupboards came to have the organisation that they do, sorting through use is a possible mechanism.[4]

4. The cupboards and drawers of Belinda's kitchen were similarly organised, with frequently accessed items stored at the bottom of high hanging cupboards (and at the top of lower storage space), and common items in drawers stored towards the front, and less frequently used items towards the back. In Belinda's case, she and her husband purposefully arranged their kitchen in this way when they first moved in. However, the organisational scheme was, they told me, a carry over from their previous kitchen and the question remains how that scheme originated.

Another emergent form of sorting can be found in several of the participants' use of their plate racks. Implements that are routinely required are likely to have been recently used and cleaned. The plate rack is therefore a probable place to find commonly used items, such as measuring cups and knives. Four of the participants in the study (Lisa, Annabel, Amanda and Robert) made frequent trips to the plate rack to fetch objects that were then used during preparation and cooking of their meals (Henry's and Belinda's kitchens lacked plate racks, whilst Benny, Elisabeth, John and Marcus didn't use theirs in the sessions filmed). Amanda had this to say about her use of the rack:

> Amanda *Those things that I use every day like cereal bowls for example and, the cup or the glass you drink juice from in the morning, everything, it's there so you don't have to, like the cutlery, the things you use most all stand there*

The plate rack is also conveniently placed in most kitchens. The processes of using and cleaning kitchen implements, therefore, sorts out, and makes readily available the most frequently used implements, making them easy to find and use.

Taking items from the rack also conserves effort. As items are taken down, less items remain to be put away when the rack is emptied; the rack is cleared partly as a by-product of the activity of cooking. Some items may never be put away, but lead an existence either in use or drying on the rack. This was the case, for instance, with the less expensive of Amanda's kitchen knives.

> Amanda *They work and they are always standing there, in the plate rack. Doesn't matter if you dry them or not* (which was apparently an issue with her more expensive knives)

Another advantage of the plate rack is that its contents, unlike those of the items still in drawers and kitchen cupboards, are visible and easily identifiable. Visibility is many times a desirable property as it can often replace the need to recall the location of some object (cf. Norman, 1988). Elisabeth, for instance, bought a magnetic strip to hang her kitchen knives on:

> Elisabeth *They* (i.e. the knives) *come in one of those blocks. Actually it's better to have them like this because then you can actually see what kind of knife it is*

The strip was fixed to the wall and carried five evenly spaced knives and a knife grinder. Placed in the wooden block that they came in, only the handles were visible. As the handles all looked pretty much the same, with only minimal variations in shape and size, they were hard to tell apart without actually withdrawing a knife from the block.

Most of the other participants stored their knives in a drawer, except for Amanda (who kept her cheap knives on the dish rack), and the two participants who had their children present during the session (John and Lisa), who kept theirs in pots on their work benches. As both John and Lisa explained:

> John [Turns to the pot on the work bench containing knives and other implements]
> *Since we have kids*
> [Grabs a knife with two fingers and jiggles it in the pot]
> *all sharp things are here. So that's thought through, it should be here. Stuff that's dangerous for the kids*

> Lisa [Gestures to the pot on her work bench]
> *I have all the dangerous stuff here, away from my son*

Regardless of what might be optimal, from a cognitive perspective, other considerations will sometimes take precedence over how things are stored. In the two cases above the issue of safety governed storage of certain items. Below is another kind of example, taken from Amanda's session:

Amanda [Takes a wooden pepper mill from the shelf above her stove and gives it a couple of twists over the sauce]
And salt
[Walks over to a cupboard. Opens cupboard and takes out a fairly large plastic brand name salt cellar]
It's so ugly it has to be kept here

Rather than letting an unattractive plastic saltcellar ruin an otherwise pastoral shelf placed above her stove (which contained other condiments), Amanda had relegated the saltcellar to the inside of a cupboard. For her, as for most of us, aesthetic preferences are sometimes more important than matters of usability. The kitchen is also a place with it's own aesthetic, something that several people showed a sensitivity to. Henry, for example, had a ceramic jar next to the stove with mostly wooden implements, as well as a plastic and steel potato masher. When I asked him about it he told me:

Henry *All the things of wood*
[Picks up a potato masher]
This shouldn't really be here, for aesthetic reasons. It's really meant for wood things because they look nice
[Touches a wooden spoon with the tips of two fingers]
but next to the microwave and my filthy stove it doesn't look so good anyway, so I don't care
[Replaces the potatoes masher]
For the same reason I like to hide things back here
[wags a finger by the side of the microwave oven near the wall]

Conclusion

The observations presented in this paper are a selection of all the things that I though myself see when I was filming the participants cooking, or when I analysed the video data collected. There were many things that did not make it into the present paper. Observations which I have omitted are ones where there were too many competing explanations for what I observed happening. Providing sufficient warrants for interpretations of data of this nature is hard. There will be interpretations that I have provided that the studious reader will find doubtful or unconvincing. For this reason I have tried to include as much transcript as necessary to make it possible for another person to reach a different conclusion to mine. In some cases the result is a long transcript and a minor point made. However, I prefer to sacrifice some style and give my reader a fair chance of reaching different conclusions. That being said, I hope that the reader will subscribe to at least some of my proffered interpretations.

This study has but skimmed the surface and I think there is much more that could be investigated. The present study has raised a number of issues that would be worthwhile to explore further. One such issue is the use of clocks in a session of cooking. I have suggested that clocks may be consulted and used for purposes other than the that of "taking time" and I also offered a number of alternative hypotheses that could be tested. An explorative study that focuses on time and clock use would be a suitable starting point.

Furthermore, since cooking is such a frequently recurring activity there is time and incentive for people to hit upon, learn or invent new tricks and techniques, and to change the organisation of their kitchens. The kitchen promises to be a good place to find special solutions tailored to problems and tasks that are frequently encountered. In the present study I visited each participant and filmed them cooking only once. It would be interesting to revisit the same group of people to observe them cooking on several occasions, to observe them cooking familiar recipes as well as novel ones, and to see them cooking various kinds of meals.

Cooking is a rich domain: it takes place in a highly structured environment and involves a number of tools and implements; there are complexities of timing and co-ordination, but at the same time there is a great deal of flexibility in how constraints are handled. We may be tempted to think of cooking as a simple activity, but that is an illusion born of our familiarity and prolonged experience with cooking and food preparation: from having stood by the stove as children, watching our parents make pancakes, to whipping up soufflés and other marvels as adults.

Acknowledgements

My sincere thanks to the ten people who took part in this study for allowing me into their homes, and to the participants of the Lund University Cognitive Science seminar for discussing a previous draft of this paper. Special thanks to Peter Gärdenfors, Lars Hall and Peter Kitzing for some especially helpful comments and, as always, to Annika Wallin for meticulous and exacting feedback. All blunders and confabulations are, needless to say, my own. I am also very grateful to the Department of Philosophy at Lund University for providing financial support as I was completing this text.

PAPER TWO

Cognitive Task Transformations

ABSTRACT: **The difficulty facing an agent performing a familiar task is one of satisfactory performance rather than one of figuring out *how* the task is to be performed. In this paper I describe and exemplify a number of principal ways in which the physical form of tools and working environments can reduce the cognitive burden of performing routine tasks. Starting from the commonly held assumption that tools transform the structure of the tasks in which they figure, I explore some basic ways in which tools can reshape tasks to circumvent excessive demands on attention, working memory and motor control, and the need to perform the kinds of cognitive operations at which we lack proficiency. The principal task transformations that are identified include: removing superfluous actions; delegating problematic parts of a task to other people or artefacts; substituting less demanding parts of a task for parts that are cognitively difficult; rearranging a task to avoid conflict between actions; and increasing the tolerance of a task environment to sub-optimal task performance.**

For each transformation an account is given of the change to the task from the perspective of an agent and how this improves the cognitive congeniality of the task. For each transformation I also give an example drawn from the everyday experience of the reader of a particular tool that may be said to bring about or facilitate the transformation in question, and explain in what way the tool is able to achieve the transformation whilst still ensuring that task demands are met.

Introduction

In this paper I describe and exemplify a number of principal ways in which the physical form of tools and working environments can reduce the cognitive burden of performing a task.

For any given task there is usually a variety of techniques and artefacts that can be used to accomplish what are basically the same utilitarian objectives. Whatever other similarities and differences there may be between tasks, from a cognitive standpoint, all ways of carrying out a task are not equal. Some ways of performing a task are easier than others, whilst some methods put a strain on our mental resources. A common example given to illustrate this point is the comparison of mental arithmetic, with arithmetic performed with the help of a calculator. In both cases the same end is served, but the paths to that end take different routes. Using a calculator does not obviate the need for mental operations, but these are different and simpler than those undertaken by a person without such a device (Cole & Griffin, 1980; Hutchins, 1995a).

Clearly, tools may do more than just furnish us with additional possibilities for action. In addition to the proper functions and symbolic functions beloved of anthropologists (e.g. Schiffer, 1992) the design of tools and working environments shape the structure of the tasks in which they figure. A tool delimits the kinds of things that can be done with it, as well as *how* those things are done.

The question of how physical structure interacts with the cognitive processes involved in carrying out a task is of prime importance, not only in design, but also in our understanding of cognition. Unless we properly understand the interaction between thoughts and things there is the risk that we will overattribute cognitive properties to people, mistaking what are in fact the capacities of people together with tools, for the accomplishments of unaided individuals (Hutchins, 1995a). To better understand cognition we need to better understand what it is that tools do, what roles they play and in what ways they may supplement our minds.

In this paper I start from the commonly held assumption (e.g. Cole & Griffin, 1980; Hutchins, 1990; Norman, 1991) that tools transform the structure of tasks. The focus here is an exploration of the subset of cognitively benign transformations that can be brought about by artefacts, and the various means by which they help achieve these. Although I will not be offering the kind of detailed and elegant cognitive ethnography championed by Hutchins (1995a, 1995b) I will provide a more systematic analysis than is usual of the principal ways in which tools can transform hard cognitive tasks into easy ones. To illustrate these basic transformations I have selected examples from everyday experience in the hope that these will be familiar to most readers.

Before going on to describe some of the ways in which tools can help reduce the cognitive burdens of task performance I need to supply the reader with some concepts and assumptions. In the next section, therefore, I explain what it means for a task to be "cognitively congenial" and also what it means for a tool to contribute to a task's cognitive congeniality.

Cognitive congeniality

The design of working environments and tools influences and constrains how a particular task is structured. Using a specified set of material resources for a task requires an agent to perform a certain sequence of actions. Different sets of tools can structure a task in different ways, some being cognitively easier to perform and some more difficult. A hard task can sometimes be transformed into an easier task by felicitous design of the material resources provided for that task.

A task may be cognitively demanding for a variety of reasons. It may be ordered in such a way, or include certain kinds of actions, that attentional resources are over-taxed. When attention is overburdened this increases the likelihood of error and can lead to a breakdown in activity. Like attention, working memory too is a limited resource that can be similarly over-taxed in the performance of a task.

Viewing cognition as a set of limited resources echoes the approach taken in human factors and much of human–computer interaction. This takes us a bit of the way, but paints an incomplete picture of cognition. In addition to these kinds of limits there is substantial evidence documenting our differential competence at various types of cognitive tasks; we simply seem to be better at some kinds of mental operations than others. It has been noted, for instance, that we seem to be particular good at recognising patterns, modelling simple dynamics of the world, and manipulating physical objects (Rumelhart, Smolensky, McClelland and Hinton, 1986). By comparison, we are notably poor at solving certain kinds of abstract logical problems. Thus some tasks may be more difficult for us simply because they involve the kinds of cognitive processes at which we are less adept.

Kirsh (1996) calls the measure of how cognitively hospitable an environment is its cognitive congeniality. A cognitively congenial environment is one that reduces "the number and cost of mental operations needed for task success," "reduces cognitive load on working memory" and increases "the speed, accuracy or robustness of performance."

The transformations that follow detail a number of principal ways in which tools can compensate for excessive demands on attention, working memory and motor control, as well as how tools can circumvent the need for us to perform the types of cognitive operations at which we are less adept.

Designing tools and tasks that allow reliable and trouble free performance is clearly desirable, but does not cover all the possible ways in which a task environment can be enhanced. In activity there may be other problems facing an agent. One such problem being the need to work out *how* to perform an unfamiliar task. But many tasks in our everyday lives are already known to us, and the problem remaining is their satisfactory execution.

Cognitive task transformations

What follows are a number of principal ways in which a task may be restructured by the introduction of a physical artefact (or by the amendment of already present artefacts) in ways that improve the cognitive congeniality of the task.[1] To view these changes as "transformations" highlights the continuity between generations of a task (since any given generation of a task is the *transformed* version of the previous generation), and the positive connotations of the word are also congruent with the positive nature of the change.

---

1. In the present paper I will only touch on the ways in which tasks can be transformed by non-material changes, such as the acquisition of new skills and procedures (for instance, new ways of handling already present artefacts).

For each transformation I shall try to provide a number of things. I will give an account of the change to the task from the perspective of an agent (i.e. how the actions that have to be performed by the agent are rearranged or replaced by other actions). Since only some changes to a task will lead to an improvement in cognitive congeniality I will also need to show how a particular transformation results in a reduction in the demands placed on cognition. For each transformation I will also give an example of a particular tool that may be said to bring about or facilitate the transformation in question and explain in what way the tool is able to achieve the transformation (i.e. in virtue of what physical properties the transformation is accomplished). Regardless of the relative cognitive difficulty of a task there will always be task demands that need to be satisfied in order for the goals of the task to be met. A crucial question is how a certain tool is able to change the structure of a task, in ways that improve the cognitive congeniality of the task, whilst still satisfying task demands. The various aspects adumbrated above will hopefully become clearer as we progress with examples and diagrams.

Although each transformation will be described as if we were deliberately setting out to improve cognitive congeniality they may also be brought about by processes lacking such intent. In addition to conscious design decisions there are non-deliberate ways in which a task may change. For instance, there may be non-intended, but fortuitous, side effects of other decisions, or accidents, that change the structure of a task. Even the process of executing or repeating a particular task may sometimes change that task in advantageous ways.

Elimination

If we want to improve the cognitive congeniality of a task one of the first things that might come to mind is to strip that task down to essentials, removing those parts of the task that contribute to its difficulty, but which are not strictly necessary for its successful completion. Many tasks and task environments are like this, requiring us to perform redundant actions, or make decisions, that would not be needed with another set of tools. There are many reasons why this may be so. In some cases it is of little importance whether a task is in its most reduced and streamlined form; demands on robustness and efficiency are only sometimes a top priority and many tasks exist in less than optimal form because other considerations take precedence. In some cases superfluous actions are the vestiges of an earlier form and organisation of a task, from a time when procedures, goals, or surrounding artefacts were different. In a similar way there are actions that may be demanded by an artefact that has been imported from some other task and then put to new uses.

These are some possible reasons for why a task may make unnecessary demands on cognition. These kinds of tasks can sometimes be rationalised by changing the properties of the material means employed in those task. In this way we may be able to eliminate parts of a task which require the kinds of cognitive operations we are less adept at, or which needlessly add to attentional load or demands on working memory.

Some tools are purposely made to fit many different kinds of task (take the Swiss army knife as a familiar example) and this may inadvertently add complexity to all, or some, of those tasks for which a tool is intended. In the case of the Swiss army knife you first have to find the right instrument before you can use it. In a task requiring you to switch quickly and fluently between, for instance, a screwdriver and a spanner, a separate tool of each kind will save you a number of actions (and mental operations) that would be demanded should you use the polyvalent army knife. In switching between instruments of a Swiss army knife you usually need to fold away the tool you are currently using, remember were the next tool is located (or alternatively spend time searching for it) and then protract the tool of choice. Conversely, the fixed alternatives to the Swiss army knife (a screwdriver and a spanner, in the example) are immediately identifiable since they have no hidden structure. Because they are independent tools, and the working bit or parts are fixed and immovable, they do not require additional manipulations prior to use.

This general principle – the elimination of superfluous actions – is illustrated diagramatically in figure 1. The top drawing of the figure illustrates a sequence of simple actions (labelled A–D) of some given task. These are the actions required by an agent performing the task before the task is restructured.

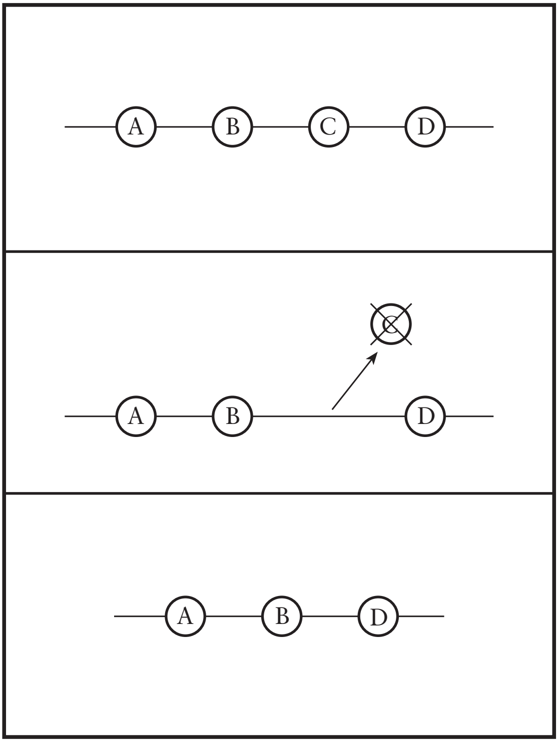

Figure 1. Elimination of a superfluous action.

The choice of representing a task by breaking it down into constituent actions is consistent with work in human–computer interaction (e.g. Diaper, 1989) as well as the anthropology of technology (e.g. Gatewood, 1985; Lemonnier, 1992). The middle drawing illustrates the task undergoing the transformation, it shows how a certain action is removed from the concerns of an agent. In some of the later diagrams the drawing also gives an indication of the relation of the actions performed to events that are outside the immediate control of the agent. The bottom drawing illustrates the resultant task facing the agent after the transformation has taken place. Since both top and bottom drawings can be inferred from the middle drawing they will be omitted in subsequent illustrations.

In cases where an action (or actions) are superfluous, in the sense that they have no essential utilitarian function in relation to the task, their disappearance from a task will not have to be compensated for. Both prior and subsequent to the transformation, task demands are met by the execution of the same set of actions, of which the eliminated action was never part. In this respect the principle of elimination differs from the other transformations. When an action that is *necessary* for a task is moved or replaced the purpose it served still has to be served in some way. In these cases, a tool that is responsible for a particular transformation also has to ensure that the actions that remain are able to satisfy task demands.

A premise of this section has been that it is possible to identify actions in a task that do not play essential utilitarian roles and which needlessly add to the cognitive burden of performing the task. Another premise has been that we might excise redundant actions from a task by suitable redesign of the material means employed in that task. There is reason to be cautious in the face of both of these assumptions. Identifying redundant actions is hard because there is always the possibility that an action does in fact have some crucial, but obscure role. For instance, a seemingly needless action may in fact belong to some other task, perhaps laying the ground for activities to come, or perhaps necessary for the co-ordination of the task with other simultaneous activities. There are also actions that assist in regulating the pace and timing of a task.

Delegation

Pruning a task can only remove some of the cognitive burden of performance and is not always an available alternative. There are usually elements of a task that are crucial and which cannot be removed if the goal of the task is still to be met.[2] When eliminating actions is not an alternative, parts of a task can sometimes be delegated to another agent or to an artefact (see figure 2). From the point of view of the agent performing the task this can have the same consequences as the principle of elimination: certain problematic parts of the task are removed and disappear from sight. From the point of view of the system comprised of both agent and task environment, the "system view" as Norman (1991) would call it, all parts of the task are still being performed, and task demands are consequently being met.

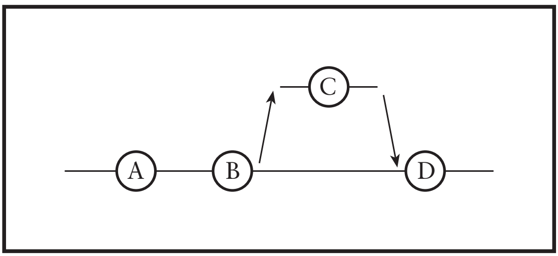

Figure 2. Delegation (without substitution).

In common parlance, when parts of a task are taken over by artefacts, we usually speak of automation or of automating a task. An everyday example of this kind of transformation is the transition to cameras with auto focus. Taking a picture once involved framing a shot, focusing the shot, and finally taking the picture. Most cameras today remove the action of focusing from the activity, but leave the rest of the task intact and unchanged. The person

2. One way in which we might identify such crucial, or "strategic operations" as Lemonnier (1992) calls them, is to compare similar technologies and techniques found in different cultures. In any technology there are certain steps that seem inescapable. Mahias' (1993) study of pottery techniques in India is an example of a comparative enterprise in this vein. Although there is great variability in the way in which pots are made throughout India there are also a number of identifiable invariants of the task. One such invariant is that pots that are fired in a kiln need first to dry out. Because of the nature of the materials employed the order of these activities cannot be changed and these basic steps can not be omitted. Caution is advised, however, since it is often these kinds of invariants which new design solutions seem sometimes to be able to change.

using a camera equipped with auto focus is left with more cognitive resources to devote to composition or timing of the shot. The act of focusing is taken over by the camera which either judges the distance to the subject (using an infrared beam or ultrasound) or compares the centre of the picture to a reference image, and moves the lens by means of a small electric motor.

Although parts of a task may be delegated there is not always a free choice of whether to delegate to another human or to an artefact. Spreading the demands on attentional resources or working memory over several people may be a good idea, but there is not always strength in numbers. Delegating parts of a task that require difficult cognitive operations may simply shift the burden from one person to another. Likewise, there are limits on what kinds of things artefacts or machines can be made to do.

Note, however, that what is being delegated are bits of the task (sub-tasks if you will) that have to be done to satisfy task demands, and that it is of no importance in what *manner* these parts of the task are subsequently performed, only that they are successfully completed. Parts of a task that are delegated to a machine do not have to be carried out in the same fashion that a human would carry them out.

The notion of delegation (cf. Akrich, 1991; Latour, 1988, 1991a, 1991b, 1999) is in some ways similar to the notion of distributed cognition (Hutchins, 1995a, 1995b), the main difference being that distribution usually designates the distributed performance of tasks that are chiefly computational in nature, whereas delegation encompasses all types of actions.

Delegation with Substitution

From the perspective of an agent performing some task, delegation can function much like elimination: parts of a task that were once difficult or problematic are removed from the responsibility of the agent. These parts are still carried out, but the agent need have no awareness that this is being done. However, delegation usually also entails some *noticeable* changes to the task. The fact that certain elements of a task have to be "handed over" to other agents or artefacts usually requires some kind of action on the part of the task performer. In most cases the "results" of these delegated sub-tasks will also have to be "returned" to the agent, in order for the activity to continue, and this will entail further changes to the structure of the task. In these sorts of cases the agent has to perform actions that enable parts of the task to be delegated, but these are actions that are less demanding than those the agent would have to perform if there was no reallocation of sub-tasks. In some forms of delegation, therefore, certain actions required by the task performer are delegated and exchanged for other, less demanding, actions (see figure 3).

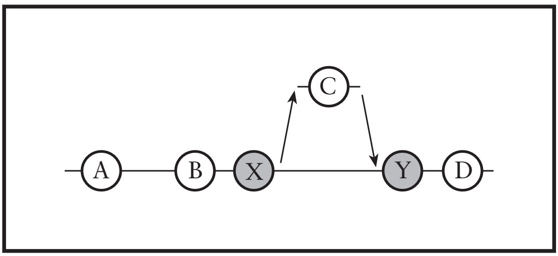

Figure 3. Delegation with substitution.

Let me give a simple example to illustrate the principle. Roadie, deck-hand and security guard are all occupations in which an occasional task is to keep a running total of the number of individuals who have passed some designated point. One way to do this, obviously, is to keep a running total in your head, but this can be surprisingly difficult. There may be distractors (a rowdy crowd, or questions asked), as well as the additional task of keeping track of which individuals have already been counted.[3] This is also a very boring task to perform and motivation may understandably falter. To facilitate the performance of this task a small and familiar device called a hand-held tally counter is sometimes used. A press on a button on the top of the counter advances an internal mechanism one step and the current total is displayed in Arabic numerals

3. Though this requirement is often eliminated by having the people that are to be counted file pass the person counting. In this way people who have already been counted are removed from view, preventing them from being counted more than once.

in a small window. From the perspective of the task performer the act of remembering the current total of individuals counted has been replaced by the actions of pressing a button and consulting a numerical display. The continuous memory for a sum has been delegated to the mechanism. It is the mechanism's ability to retain a configurative state that allows it to serve this role.

Making the world remember in our stead is a pervasive strategy. We routinely write things down on paper, of course, but we can also encode information in spatial arrangements (Kirsh, 1995b) and by using meaningful categories of objects (Beach, 1988). These mnemonic strategies all change what are basic memory tasks into encoding–decoding tasks. The task environment has properties which make this possible (among those being: sufficient plasticity to permit a change in state, a sufficient longevity of the new state, and the visibility of this state. cf. Latour, 1986) and therefore allow us to substitute the kinds of cognitive operations that have to be performed. From a functional standpoint the role played by the environment here is very close to biological memory (cf. Clark and Chalmers, 1998).

Delegation of sub-tasks and substitution of actions commonly go hand in hand. The examples given here are just a few of all the possible permutations that this way of transforming a task allows. Of the transformations detailed in this paper this is probably the task transformation that has the greatest generality and potential.

Substitution

By substituting an easy action for a cognitively demanding one (demanding because of its nature or because of its relative position in an activity) the cognitive congeniality of a task can sometimes be improved. We saw above how the substitution of an action is commonly an integral part of delegation, but the principle of substitution can also be employed on its own (see figure 4). When the substitution of an action is unaccompanied by delegation, the new action has to ensure that task demands are being met. To put it in another way, the physical artefact that has been modified or introduced has to enable some new action to satisfy task demands, but to do so without actually taking over any aspect of task performance. This distinction between delegation accompanied by substitution, and substitution on its own is not altogether easy to draw. Some examples will help make things clearer.

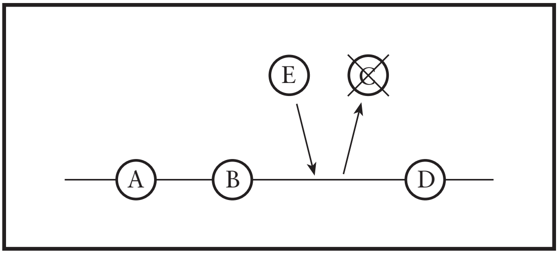

Figure 4. Substitution (without delegation).

Remembering arbitrary sequences of numbers is a component of such activities as logging in to a computer and withdrawing money from an ATM. Remembering and recalling strings of meaningless numbers is, however, not something we are generally good at. A task in which we are required to input a string of such numbers can therefore be simplified by exchanging the representational format of the thing to be remembered for something more in tune with our cognitive makeup. Something we are quite good at remembering is people's faces and this is a fact exploited by the Passface™ authentication system produced by ID Arts. As the name of the system suggests numerical passcodes have been exchanges for, so called, "passfaces". Using this system the task of remembering numbers and inputting them (which is cognitively demanding), is exchanged for the task of recognising a number of known faces from a larger array of distractors (cognitively less demanding).[4]

An example of a tool that permits the use of cognitive abilities we are especially good at is the slide rule. This is an artefact that transforms a complex computational task into one of relatively simple scale-alignment (Hutchins, 1995a). Operations requiring

4. Note that the system not only changes the representational format of the-thing-to-be-remembered, but also changes a task requiring recall into one requiring recognition (at which we are far more proficient). For more information on the Passface™ system pay a visit to www.realuser.com.

relatively simple physical co-ordination are substituted for more complex mental operations. Whether this is to be taken as an example of substitution (pure), or of substitution accompanied by delegation, hinges somewhat on what we understand the slide rule to be doing. A slide rule maps numerical operations onto simple movements thus allowing physical manipulations to execute mathematical computations.

In addition to the option of substituting a less demanding action for one that is more demanding, an action can also be exchanged for several. A part of a task that makes great demands on motor control or attention can sometimes be divided up into several smaller actions where each is within acceptable resource limits (see figure 5).

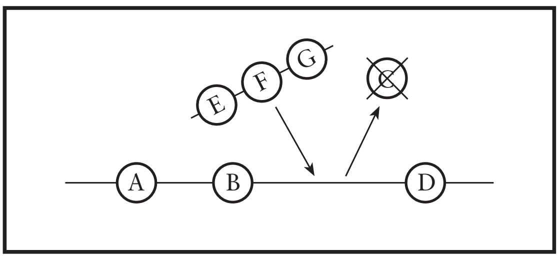

Figure 5. Substituting several actions for one.

Parts of a task that require skilled physical manipulation can sometimes be exchanged for several actions, each requiring simpler kinds of manipulation. If the complexity of a manipulation is due to the number of degrees of freedom that have to be simultaneously controlled, suitable artefact design can make it possible to control these degrees of freedom sequentially. This is the strategy used when planing a piece of wood at a workbench. The piece being worked, and the tools used, are all set up in such a way that we only have to control one degree of freedom at a time.[5] The

5. Another way of handling difficult motor actions is to substitute manipulations that are already part of the agent's repertoire for the manipulations previously required. These manipulations may still be complex, but are ones at which the agent has had sufficient training.

skilled action required without the use of the bench is exchanged for sequences of simple tool management (cf. Blackaby, 1986).

When attention is over-taxed a possible remedy is to shift some of the demands on attention over to a neighbouring channel. A task requiring complex visual monitoring can sometimes be divided into a task requiring visual *and* auditory monitoring. The strategy is to subdivide a problematic aspect of the task and substitute simpler activities.

Bærentsen (1989) gives an account of the early history of firearms. He notes how the physical design of the earliest guns were the cause of conflicting demands on visual attention when the guns were to be aimed and fired. Firing a gun was achieved by placing a burning fuse through a small hole on the top of the gun. Visual attention was required in order to find the hole and to monitor the motor actions of fitting the fuse into it. Unfortunately this also left little visual attention available for aiming. In later generations of the gun the fuse was mounted on the end of a mechanical arm that was fitted to the side of the gun. Firing the gun could now be achieved by pulling on the mechanical arm, which caused the fuse to dip into the hole and ignite the gunpowder. In this case part of the attentional control demanded by the task was shifted over into a neighbouring channel (de Léon, 1999).

Rearrangement

The difficulty of a task is sometimes the result of how parts of a task are ordered. A particular sequence, for example, may place difficult parts adjacent to one another, whilst another sequencing of actions can be more benign. A difficulty may be in moving from a particular action to the next, or excessive demands on memory and attention may simply accrue in non-beneficial ways. In these sorts of cases a possible remedy can sometimes be to reorder the sequence in which the steps of a task are performed (see figure 6). For some tasks this will not require any change in the tools and materials used in the task, but rather a change in procedure. In other tasks a reorganisation will also require a redesign of surrounding artefacts.

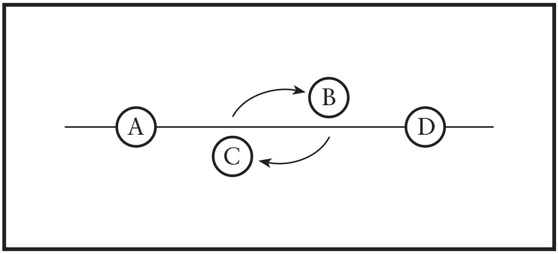

Figure 6. Rearrangement of actions.

For an example of how memory demands may be linked to task organisation consider two ways in which someone working in a newsstand might assemble an order for multiple items. One strategy is to retain the full order in memory and then to fetch the items requested. An alternative strategy would be to assemble the order as its being made, picking out each item as its named. The two strategies involve basically the same actions and operations, of remembering, locating and fetching a number of items. The second strategy, however, requires less to be held in memory at any given moment. For the second strategy to be available to an agent the newsstand has to be physically constructed and organised in a way that places various item within easy reach.

One way of managing excessive attentional demands is to permit the performance of certain actions well ahead of the time at which they are needed, what Hutchins (1995a) calls precomputation. This strategy requires that actions can be performed ahead of time and "saved" in a form that enables them to be reanimated at some later point when they are needed.

Tolerance

The last of the transformations takes another kind of tack. Rather than alleviating the demands placed on the cognition of an agent performing a task, we can sometimes change the *consequences* of sub-optimal performance. If we know that certain kinds of errors are likely to result we may be able to construct a task environment that is tolerant of those errors and which permits less than perfect performance to satisfy task demands. If we know, for instance, that attentional or working memory resources are likely to be over-taxed we can make predictions regarding the kinds of errors or mistakes that people will be prone to make. Together with an understanding of the goals associated with the task, less than perfect performance can then be charitably "interpreted" to yield results that are likely to be those that were intended. From the perspective of an agent the task still has the same structure as it had prior to its transformation: it is still the same actions that have to be performed. The difference is that the task will show a greater tolerance of how well the agent performs the requisite actions. In the figure illustrating this principle (figure 7) we see how performing any of three variants of a certain action will bring an agent through the task.

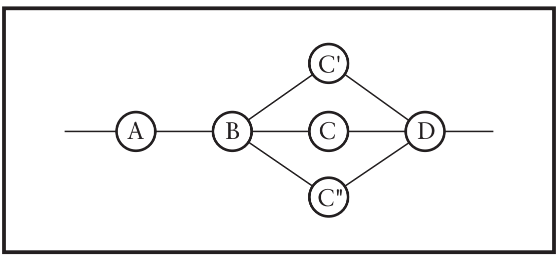

Figure 7. Tolerance.

Pouring liquids with the help of a funnel is an example of this design strategy. The funnel permits a greater range of "acceptable" motor actions to result in water being poured into a receptacle (rather than down the side of it). The same results can be attained with a lesser degree of control.[6]

A needle-threader is another elegant, but less well known, artefact that makes use of this strategy. The implement consists of a small metal disc, large enough to be grasped between thumb and

6. The funnel introduces a number of additional actions to the task (e.g. those required in appropriate placement of the object). The strategy is, however, distinct from the other principal transformations. You can imagine, if you like, a restaurant replacing all its narrow necked carafes with wide necked ones.

forefinger, and has a small wire loop protruding from its edge (see figure 8). To thread a needle using this artefact you first push the loop of the threader through the eye of a needle. The loop is larger than the hole in the needle, but collapses as its pulled through. In addition, the far side of the loop of the needle-threader is pointed to facilitate its passage through the needle (the stiffness of the wire as well as the point makes it easier to put through a needle than a flimsy and snarled end of thread). Once the loop is through the needle, thread is put through the loop (the loop being larger than the eye of the needle makes it easier to thread). The loop is then extracted from the needle, pulling the tread with it and through the hole.

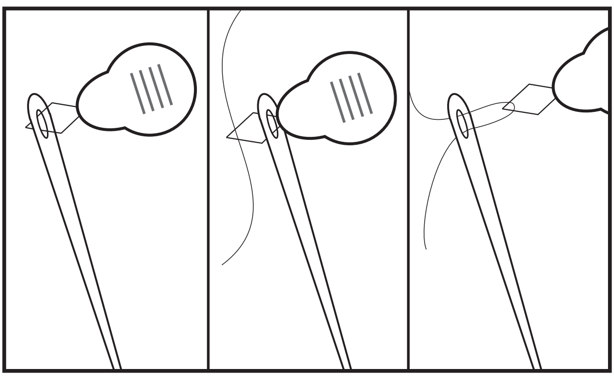

Figure 8. How to use a needle-threader.

Other material means of improving cognitive congeniality

In the above I have explored some basic ways in which the material means provided for a task are able to influence or determine a task's organisation. The approach is one of fitting tools and tasks to minds. This approach does not encompass all the different ways in which physical structure might aid a task performer. The focus is on facilitating the performance of a task which is already known to the agent, where the main problem for the agent is not one of figuring out what the task is or how it is to be executed, but performing it in a satisfactory manner (Hutchins, 1990). Obviously this does not cover all types of tasks, but a large part of everyday activity is like this. It is also the case that tasks for which special artefacts have been developed or adapted have to be routine tasks: it is only when a task is repeated sufficiently often that special artefacts are warranted and where there will have been time for them to develop.[7]

Another limitation of this approach is that I have considered the work of a single individual and not the work of a concerted group of individuals. A place to turn for accounts of how artefacts support the work of several co-operating individuals is the field of computer supported co-operative work, and of course the work of Hutchins (1995a, 1995b).

The problem for an agent of figuring out how to perform a task, an issue I self-confessedly do not address, is also something which can be aided by appropriate design of material structures. Norman (1986, 1988) stresses the potential of design to communicate a conducive conceptual model to a person using a device. Interestingly, the mental model communicated need not necessarily be a veridical reflection of the device and task in order to be productive (Norman, 1988).

Another way of conceptualising the problem facing an agent is to see it as one of proper action selection, of knowing what to do next at any given moment (Kirsh, 1995b). When understanding of a task is lacking and there is no prior representations of the task to consult, the task environment can be sometimes be augmented in ways that can facilitate appropriate action selection.[8]

---

7. See, for instance, Keller and Keller's (1996) study of modern blacksmiths. The smiths will usually use what they call tools-of-risk, tools which require high levels of skill to use and where the likelihood of mistakes are high. But when orders for multiple copies of an object are received, they commonly construct so called tools-of-certainty. These are jigs and other one-off tools, specially made for the occasion, which make repetitive construction more of a routine affair.
8. Kirsh's (1995b) paper details a number of different ways in which the spatial configurations of objects can aid action selection, but the principles can be translated into prescriptions for artefact design without too much strain.

Conclusion

A number of basic ways in which tools can improve the cognitive congeniality of a task have been presented. These do not tell the complete story of how cognition relates to material structures – thinking to things – but serves, together with the work of others, as a beginning. The relevance of this kind of enterprise should be apparent. First, it may be of use in the endeavour to design better tools and task environments. We should always be cautious in assuming the relevance of theory to the practice of design, but perhaps a framework, like the one presented here, can be employed in task analysis, to highlight difficulties and suggest alternatives. In the end, of course, it will be people who decide what works and what doesn't. And this brings me to another area in which the task transformations may be of some relevance. The area I am thinking of is the study of artefact evolution. Although the development of artefacts has long been the domain of archaeologists and historians of technology there are many ways in which cognitive science might participate. On the one hand, we may be able to explain certain innovations and changes to artefacts as moves that improve the cognitive congeniality of particular tasks. This will be one more tool to add to the symbolic, functional, economic, historical, cultural and social explanations. On the other hand, we can study how artefacts have evolved, and continue to evolve, in order to learn more about what kinds of cognitive tasks, we as practitioners, try to avoid, and what kinds of tasks we continually gravitate towards. By learning more about how task and tools evolve we may learn more about cognition.

Acknowledgements

I would like to thank David Kirsh, Henrik Gedenryd, Annika Wallin and an anonymous reviewer for some very helpful comments and suggestions. I am also grateful to the participants of the Lund University Cognitive Science seminar for their input on an earlier draft of the paper. This work was supported by the Swedish Foundation for Strategic Research.

PAPER THREE

The Cognitive Biographies of Things

ABSTRACT: In this paper I argue that there is much to gain by constructing "cognitive biographies of things". These are accounts that detail the history of an artefact and its use, focus on the physical changes undergone by the artefact over time, and draw out the cognitive corollaries of those changes.

There are a number of reason for why we might want to construct such biographies. If we subscribe to a situated view of cognition, then at least some of our cognitive achievement stem from the physical structures around us. To ignore the genesis, development and appropriation of cognitively significant structures would then be to paint an incomplete and misleading picture of cognition.

Even without commitment to situated cognition, constructing such use-histories may be required if we are to discern the cognitive roles currently being played by some artefact (which may be of interest, for instance, in the process of designing artefacts). By overlaying succeeding phases of an activity with preceding ones, areas of possible cognitive significance can be highlighted and explored.

In this paper I take a shot at constructing a cognitive biography of a large spice shelf that I encountered whilst conducting a study of people cooking. The biography is based on video of a session of cooking as well as a number of interviews. The resulting account spans a period of roughly 30 and details the mutual influence between cognition, activity and changing physical structures. In constructing this biography I hope to show some of the possible difficulties of doing so, as well as some of the benefits.

Introduction

We are currently seeing an increasing awareness in the cognitive and mind sciences of the importance of physical structure for cognition. The old view of cognition as something that takes place only in the head has been replaced (or at least tempered) by a view that recognises the roles played by physical and social structures. The environment, it turns out, is not just an arena for action – a playground for problem solving and plan constructing minds – but is intimately implied in many, if not all, cognitive processes.

There are several different ways of advancing this claim, but most would agree that physical structures in the world can, at least, act: as an extension of memory (Hutchins, 1995b; Beach, 1988; Norman, 1988), to simplify choice, perception and internal computation (Kirsh, 1995b; Clark, 1997), to constrain and even determine cognitive behaviour (Zhang & Norman, 1994), and to transform tasks in ways that better harmonise with our cognitive competencies (Hutchins, 1990, 1995a; Norman, 1991).

Once this basic idea has been accepted it should be natural to ask *how* these kinds of physical structures come about. After all, if physical structure can be an intimate and integral part of cognition, more so than previously recognised, then asking questions about the growth, development and appropriation of these structures should be as natural as asking age old questions about learning and development. In fact, both learning and development need to be understood afresh in the light of these emerging insights. The genesis, evolution and adjustment to cognitively significant structure ought to be viewed as an essential aspect of most, if not all, of our cognitive achievements.

Most of the authors cited above *do* acknowledge that there are interesting processes responsible for the build-up of cognitively significant physical structure, but these processes then figure to a negligible extent in their accounts.

Hutchins (1995a), for instance, gives truncated histories of the astrolabe and the compass rose (both ancient navigational instruments that significantly transformed the cognitive task of ship navigation), but then simply concludes that practice can be "crystallised" into things, without discussing the process of crystallisation itself.

Kirsh (1995b) acknowledges that the interaction of agent and environment can be studied along different time scales, and does an admirable job of looking at medium and short term mechanisms of how people set up their workplaces for particular tasks. However, the issue of how the workplaces themselves evolve is not addressed.

Bærentsen's (1989) work on the evolution of the rifle is an inspirational exception which explicitly deals with the interaction of artefacts on the cognitive demands of task performance, and the influence of cognition on artefact development. Although a bold and innovative attempt, Bærentsen's analysis relies on an unanalysed notion of cognitive processes being built into things (for a critique see de Léon, 1999).

Activity theory (the tradition in which Bærentsen's paper is written) places great emphasis on the historical and cultural foundation of thought and artefacts and would therefore seem to be an ideal place to find the kind of analysis sought for. The activity theoretical concepts of *externalisation* and *historicity* also seem to capture the concerns discussed. However, as Engeström (1999) has noted, there seems to be a general paucity of work in activity theory on these very topics.

Naturally, a part of the reason for this neglect is connected with the difficulty of reconstructing what are primarily historical processes. Unless we limit our interest in the ways in which artefacts and practice co-evolve to very short time scales (for a nice study in this vein see Agre & Shrager, 1990) we have to choose between longitudinal studies and historical reconstructions. Longitudinal studies are in many ways ideal, but the method demands great effort without guaranteeing results (even if we were to chose our sites intelligently). Reconstructions (regardless of the principles that are used to govern the reconstruction) will simply involve too much speculation for some people's taste.

Despite these concerns we should not be too quick in dismissing these lines of inquiry as each of the approaches has its potential benefits.

In this paper I have opted for an exploration of the reconstructive approach. Of the options presented it is the approach that, at first sight, seems beset by the most worrisome methodological concerns. However, it is also an approach that has some possible advantages and some unique characteristics. In contrast to more constrained studies of phenomena taking place on a short time-scale it allows for investigation of real-life events spanning changes over long time periods. Although less controlled and more speculative than a longitudinal study might be it permits us to explore sites where the occurrence of artefactual change has already been established.

In this paper I will take a shot at constructing what I like to call a cognitive biography, tracing the life-history of a particular artefact and its use over a period of roughly 30 years, and detailing the mutual influence between cognition, activity and changing physical structures. In contrast to reconstructions of events taking place over several life times (cf. de Léon, 1999) the present time scale makes ethnographic methods and structured interviews part of the available methodological arsenal.

The artefact in question is an unusually large spice shelf that I encountered whilst conducting a cognitive ethnography (Hutchins, 1995a, 1995b; Lave, Murtaugh & de la Rocha, 1984) of people cooking in their home kitchens (see de Léon, 2003b). Each of the participants of the cooking study were video filmed whilst preparing dinner and later interviewed about the organisation of their kitchens and about the origins of their tools and cooking practices. One of the participants of the study was Robert, a man in his mid-fifties. It is in his home that the shelf described in this paper resides. Based on the interview conducted at the time of the study I have attempted a reconstruction of the genesis and evolution of the shelf and concocted a credible story of the underlying factors behind the various changes to the shelf, as well as their probable cognitive consequences. The reconstruction was continuously checked against the tape I had of Robert cooking and a number of supplementary interviews. The shelf did not always appear as it does now. The collection of spices have been stored in a manner of different ways and the actual shelf makes its appearance fairly late in the story. Although I speak of the evolution of "the shelf" it is really the history of a constellation of artefacts and practices.

This paper is thus an experiment in reconstructive cognitive biography. The result is an unusual hybrid: on some occasions I use data from the case study to make particular claims, at other times I introduce extraneous theories and observations to bear on the case in question and to explain my observations. I hope that the attempt might give some indication of what a cognitive biography might look like, and what sorts of things we might learn by constructing them.

A COGNITIVE BIOGRAPHY OF A SHELF OF SPICES

First I will give a brief description of the shelf and then outline some ways in which its current structure and organisation supports Robert's cognitive activity whilst cooking. This is followed by a reconstruction of the shelf's history.

A Brief Description of the Shelf

Most of the spices in Robert's kitchen are kept on a tall shelf fixed to a wall, a few steps from the stove and workbench were the main activities of cooking take place (see picture 1). The shelf consists of two prefabricated units bought at IKEA (a ubiquitous Swedish furniture store) that have been placed one above the other and glazed in blue to match the other fixtures in the room

Each plane of the shelf is just deep enough to accommodate one spice jar and wide enough to accommodate a row of about ten jars. All jars have been labelled with embossed plastic strips and are neatly aligned along the shelving; almost all are of identical size. Row upon row of yellow plastic lids divulge their origins as reused Coleman's mustard jars. Dispersed among these are a couple of tins, a few brand name spice jars and two pepper mills of disparate design. Through the clear glass of the jars various dry powders, seeds, flakes and roots can be seen, their colours ranging from beige to brown.

In the narrow space between the two units that make up the shelf – too narrow to accommodate an additional row of jars – lies a bulldog clip, some rubber bands, a small pile of paper twist ties, and a paper packet of black pepper. Along one side of the shelf hangs a decimated garlic braid.

In all, almost exactly one hundred jars and containers are kept here. By any measure this constitutes an unusually large and impressive collection.

Picture 1. The shelf (to the left of the door), sink and plate rack (far left of picture) and corner of fridge (far right side of picture). Part of the dining room can be seen through the kitchen door.

The Organisation of the Shelf

Robert's cooking encompasses Western as well as more exotic cuisines. About a third of the spices found on the shelf (the lower three planes) are endemic to French and Italian cooking, the remaining shelves being devoted to spices used in Indian, Middle Eastern, Chinese, Japanese, Indonesian and Thai dishes.

The spices are organised into categories according to a number different principles (see figure 1 for the basic categorisation of the shelf).[9] The top shelf, for instance, contains various hot spices, such as chilli and cayenne. Here it is taste and function that deter-

9. The categories in figure 1, and in the text that follows, are Robert's own and were taken from a picture of the shelf which he had annotated. In the interview Robert also explicated some of the uses to which he put the spices.

mines placement. Another section contains different kinds of dry leaves used in South and East Asian cooking, and is thus loosely organised by the *form* in which the spices are found. A less obvious category is the group of spices placed together because of their modest application, being dubbed "by-the-pinch spices" by Robert; the principle governing this particular categorisation being the manner in which they are applied. There are also several sections in which the spices are grouped together because they are used together in a particular style of cooking. There are prominent sections with spices used in, for instance, Indian and Chinese cooking.

In those cases in which a particular spice is found in more than one form (e.g. whole and ground) these are placed adjacent to one another.

The bottom three shelves are home to more familiar spices used in Western cooking. This is the largest regional section and has been arrayed alphabetically. The spices kept here are shared by Robert and his wife. Since his wife lacks his penchant for spicy and exotic food this is the only part of the shelf that she ever uses. As she is also substantially shorter than her husband the placement of the spices, in this mutually accessible region, is particularly felicitous. The reason Robert gives for the alphabetical ordering of these shelves is that it was the only obvious categorisation to present itself that would serve two people. It is instructive to note that the only region of the shelf that is used by more than one person relies on a culturally conventional system of categorisation. Alphabetisation ensures a clear, mutually intelligible, and maintainable order.

There is also a small section of spices that are "on their way out". According to Robert these are spices that will be sacrificed as space is required. Some of these are spices that Robert once bought out of curiosity, or that he has finished experimenting with, or for which he has recently found better alternatives (for instance, access to fresh spices that were only previously available in desiccated form).

Some Ways in Which the Shelf Currently Supports Cognition

Since each plane of the shelf is only deep enough for one jar, almost all the jars can be seen at any given time; of the roughly one hundred jars that stand here only a handful are blocked, or partially obscured, from sight. The jars are labelled and their contents can be clearly seen through the glass: label and content mutually contributing to the ease of identification of a particular spice. We might think of the shelf as a kind of conceptual model, a physical structure embodying the basic spice combinations of some of the most common, Asian and Western dishes, as well as certain aspects of Robert's cooking practice and his personal way of conceptualising cooking.

We should be cautious, however, lest we view the shelf simply as reflection, projection, or externalisation of Robert's inner representations of spices and cooking. For one, the genesis of the shelf belies such an interpretation. As we shall see in the following, there are several determinants of the shelf's structure and organisation that are incidental, rather than intentional. Moreover, to whatever extent (and in whatever form) the structure and content of the shelf is actually represented by Robert, this will have been repeatedly shaped by the presence of the shelf itself.

The visible thematic spice groups arguably function to support Robert's memory in a number of different ways. Assuming that he knows what spice he is looking for, a problem then facing him is locating the appropriate jar on the shelf. Instead of having to scan the entire shelf for the spice in question, the thematic groups focus his search to a particular area (obviating the need for knowing the exact location of every spice). For example, just knowing that a spice is, or is not, used in Western cooking removes a large part of the shelf from consideration. That the organisation of the shelf is consonant with his own idiosyncratic categorisations, conceptualisations and habits of cooking also makes the regions more easy to locate.

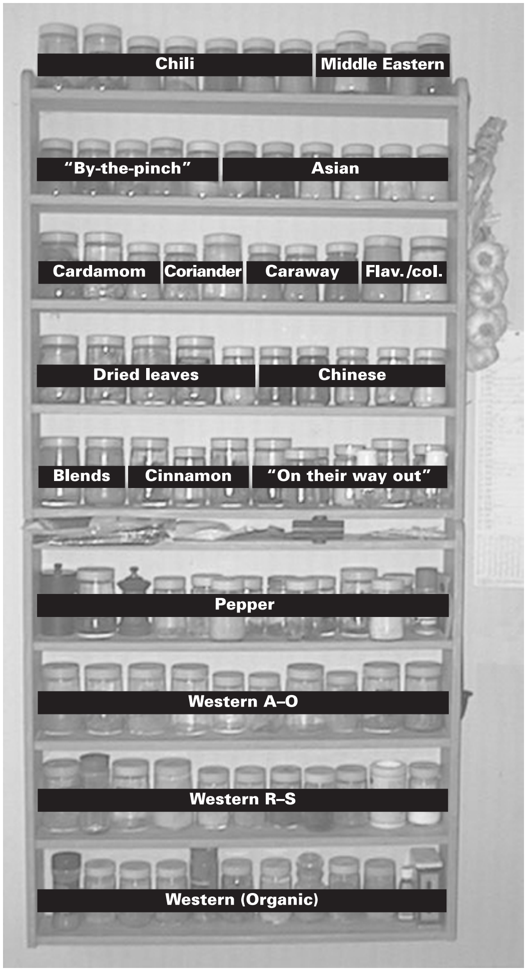

Figure 1. Basic categorisation of the shelf. The fastidious reader might note a slight discrepancy between the labels in the picture and the description of the categories given in the text. The figure is based on a picture taken a year after the original study and in the interim the organisation of the shelf changed somewhat (more on this below).

Even the relationship between the form of his body and the position and organisation of the shelf comes into play here. The spices that are most frequently used (the Indian spices, according to Robert) are within easy reach, whereas slightly less frequently used spices (such as the Middle Eastern spices) require him to stretch upwards, or (in the case of Western cooking) to bend down. The placement of the various groups of spices is not only a question of physical effort and comfort, but ensures that the customary position of Robert's body when facing the shelf presents him with the most regularly used sections of the shelf. The same holds true for Robert's wife, whose length constrains which shelves she has easy access to. This is a very straightforward and clear example of the role of embodiment for cognition (for a critical review of the concept of embodiment see Clark, 1999a).

These, then, are some ways in which the shelf can support the task of locating a particular spice. When cooking a dish from memory an occasional problem is actually remembering which spices to look for in the first place. In such a case, all Robert needs to know is what *kind* of spices he is searching for. Looking at the appropriate area of the shelf he need then only recognise the spices required – a far easier task than recalling them (for a primer on recognition and recall see Baddeley, 1997).

The co-location of spices that are commonly used together also serves as a reminder, throughout the cooking session, of spices still to be applied. Anytime a spice jar is replaced on its shelf, neighbouring jars can jog Robert's memory.

There is also an interesting structural feature of the shelf that supports the replacement of spice jars. Since the spices are stored on the shelf one jar deep, and each plane is full to the breadth, the removal of any jar leaves a clearly visible gap. This gap can then function as a perceptual place marker. Again, this saves Robert from having to remember the exact locations of where particular spices are stored: he has only to look for a break in the array to know where to replace a jar. Of course, should Robert pick out several jars together, he would be left with the problem of pairing the correct jars with the appropriate gaps when time comes to replace them. This is not an insurmountable problem and would still be easier than having to remember the exact locations. However, as seen on the video of him cooking, Robert only picks out and uses one jar at a time, replacing it before picking out the next one, thus ensuring that there is no confusion as to which jars belong in which gaps. There are several ways in which the shelf could potentially be used, but significantly, Robert has settled on a strategy that *permits* the physical structure to simplify the cognitive demands of correctly replacing used jars. It is well to point out that it is the *combination* of techniques, procedures and/or habits, with particular artefacts and task environments, in relation to specific tasks, that determines the cognitive congeniality[10] of an activity.

THE HISTORY OF THE SHELF

If we want to understand how the shelf came to have its present structure and use, we must go back in time to the late sixties, to the spice shelf in one of Robert's previous homes, and then trace the development of the shelf back to the present. And this is what I will attempt next, reconstructing the genealogy of the shelf with Robert's help.[11]

In the first kitchen that Robert remembers having (in a flat in Stockholm, Sweden) the spices were kept – in no determinate order – on a single shelf next to a stove (see step 1 of figure 2). At that time his collection was considerably smaller, consisting of only a few conventional Western spices, and could fit snugly on a single plane. Robert had yet to develop the intense, and broad, interest in cooking that he has today, and exotic food was still something that was only occasionally sampled in one of Stockholm's few Chinese restaurants (in the late sixties Stockholm had only two, or possibly three, of these).

10. Kirsh (1996) calls the measure of how cognitively hospitable an environment is its "cognitive congeniality". A cognitively congenial environment is one that reduces "the number and cost of mental operations needed for task success," "reduces cognitive load on working memory" and increases "the speed, accuracy or robustness of performance" (Kirsh, 1996). The use of the term is extended here (for reasons outlined in the above paragraph) to cover activities as well as environments.

11. I have taken some liberties in truncating the history of the shelf and excising portions that add little, from the point of view of the reader.

Although there was no intentional order to the arrangement of the jars at that time we can speculate that there may have been some incidental grouping resulting from the handling of the jars. In the kind of arrangement described, spices that are often used together will tend to gravitate towards each other. There are several ways in which they may do this. If all, or some, of the spices used in a particular dish are taken off the shelf and put back at the end of a session of cooking, spices that often occur together will tend to end up in close proximity. Or, if spices are placed at the front of the shelf immediately after use, more commonly used spices will be found near the front, whilst spices in less demand will be gradually pushed towards the back wall. Needless to say, this kind of grouping would have aided the location of common, and even uncommon, spices. With just a few spices in the collection there would have been little incentive to organise them further.

These kinds of processes, in which the repeated performance of an activity shapes the environment in which that activity takes place, are an important source of supportive structure.[12] The spatial redistribution of artefacts is one basic mechanism which we have already encountered. In the same way that spice jars can come to be functionally grouped through their use, so may other artefacts employed in the kitchen. During my study of cooking (de Léon, 2003b), from which this case is taken, I noticed several of the participants making frequent trips to the plate rack in order to fetch common objects. Implements that are routinely required are likely to have been recently used and cleaned. The plate rack is therefore the most probable place to find them. In addition, the rack is also conveniently placed in most kitchens and its contents visible. The processes of using and cleaning kitchen implements, therefore, sorts out, and makes readily available the most frequently used implements.

Note that this particular use of the rack also conserves effort, the effort of emptying the rack and replacing the things kept there. The rack is cleared almost as a by-product of the activity of cook-

---

12. Barker (1968) calls this kind of relation between behaviour and environment "behaviour-milieu synomorphs".

ing.[13] Much of human activity is like this, with actions having multiple reasons and serving multiple simultaneous purposes (cf. Wertsch, 1998).

Basic wear and tear is another mechanism that can generate supportive structure. For example, think of a footpath kept clear through use (see Barker, 1968), or the growth of meaningful pathways between buildings (see Ribeiro, 1996). Or, consider how repeatedly opening the phone book to the same section (for instance, the pages containing your local pizzerias) can weaken the spine of the book making it easier to locate those numbers in the future.[14] In a similar manner, tools and implements that are stored in designated locations (for instance, around the walls of a workshop) can discolour, or otherwise mark, those places, thus facilitating their correct replacement.

In these examples activity results in some cognitively congenial change to the environment. The natural response to such change is compliance; the net effect is beneficial and we probably do not pay it much attention. However, it is more common for activity to have the opposite effect, creating clutter and disorder, rather than order. One possible way to respond to encroaching entropy is to actively counteract it, a strategy which Hammond (1990) calls enforcement.[15] For example, in the video of Robert cooking there is a slight lull in the session which he spends meticulously straightening his spice jars, making sure that all "misaligned" labels face to the front. As the shelf is subject to rather heavy use (in the session filmed a total of 18 different spices were taken down and replaced) this kind of upkeep becomes a natural part of activity, ensuring that the shelf can continue to function as it does. Of

13. Reading an earlier draft of this paper one of my colleagues pointed out an important exception to this way of using of the plate rack. Although she recognised the use of the rack described here, she mentioned that she will sometimes leave the rack untouched for a period of time after washing up so that the full extent of her domestic efforts might be recognised, and appreciated, by her boyfriend.
14. This is why your avant-garde books always seem to fall open at the raunchy episodes in the hands of any guest browsing your bookshelves.
15. Hammond (1990) and Hammond, Converse and Grass (1995) take enforcement to be an *active* strategy of order imposed on the environment. It is interesting to note how, in the present case, enforcement is sometimes *incidental* (recall how clearing the plate rack can be a by-product of other activities).

course, it could also be argued that the episode simply reflects an aesthetic preference or ideal, rather than maintenance of cognitively significant structure (although neither interpretation invalidates the other). The physical properties of the shelf and jars also help to keep entropy at bay. The width of the shelf, for instance, greatly limits the ways in which the jars can shift about on the shelf, and the gaps left by jars that have been removed facilitate their correct replacement.[16]

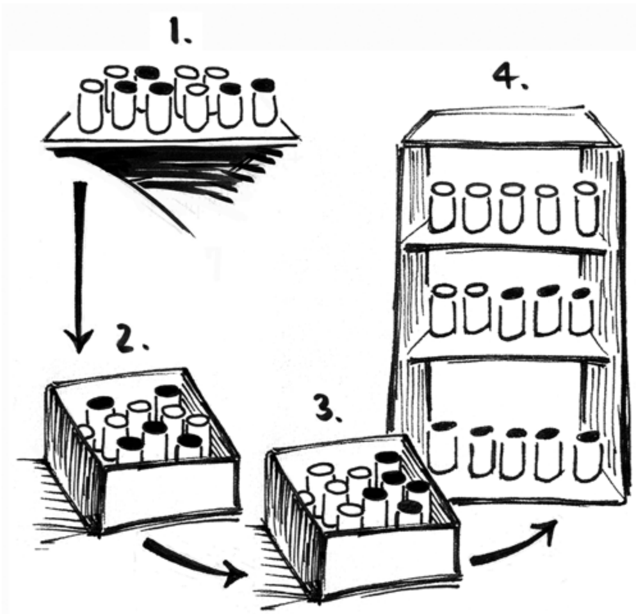

Figure 2. A schematic history of the shelf.

16. It would be interesting to analyse and compare the various artefactual and procedural means that have been devised to resist entropy. What technical solutions are there? For what kinds of task is artefactual stability (taken to mean something other than plain robustness) a desiderata? Does stability stand in the way of other functional aspects? If so, what kinds of trade-offs are there (cf. Bleed (1986) on the trade-offs between reliability and maintainability in hunting weapons)? An obvious place to start would be to look at collections and catalogues and how these are used, and how they have evolved. Or what about surgical instruments, how are these handled and cared for?

The Collection Grows

Towards the beginning of the seventies Robert starts to experiment with Chinese food in an attempt to recreate some of his favourite restaurant meals. There are, as yet, no Chinese cook books available in Swedish (and books in English are still hard to come by). However, the Swedish-Chinese Association publishes a small pamphlet that Robert procures. Through an acquaintance (a supplier to some of Stockholm's delicatessens) Robert buys exotic spices in small yellow stackable tins. Throughout the seventies Robert's spice collection grows, in concert with his steadily increasing knowledge and interest in cooking. Robert and his family move a couple of times and at some point Indonesian dishes are added to the repertoire. The small yellow tins are gradually abandoned as spices become more readily available from other sources. The spices, that are now bought by weight, are transferred to recycled Coleman's mustard jars (these are the jars that can be seen in picture 1 and figure 1).

Spices in a Box

In 1979 Robert and his family move from Sweden to England. An extended period followed in which they lived in a succession of more temporary settings: a van, a trailer, a couple of rented apartments. During this transitional period Robert kept his spices in a low box (see step 2 of figure 2).

Looking down into the open box only the lids of the spice jars could now be seen, the many identical lids exacerbating the difficulty of distinguishing one spice jar from another. To find a particular spice, Robert had to rely on his memory and/or make an educated guess. Whether or not the jar selected was the one actually sought for he would have to lift it above its neighbours to be certain of its identity. Some of the incidental ordering on the shelf may have survived the transfer to the box, but may also have been broken up. Assuming that at least some of the ordering of the spices made it into the box an incorrect guess might give a clue to whether he was searching in the right area of the box. Again, just using the box may have brought some gradual ordering to the spices.

The temporary solution, as so often is the case, turned out to be less temporary than initially expected. Eventually tiring of the impracticalities of the arrangement Robert decided to order the spices into thematic groups (see step 3 of figure 2). To locate a jar he would still have to rely on memory. Nevertheless, the organised box was an improvement over the earlier, mostly haphazard, distribution.

This is the first full and deliberate ordering of the spices undertaken. If there was already some order in place, as I have suggested, we can speculate that that order may have influenced the subsequent intentional organisation, perhaps serving as a rough guide. Any ordering that is conserved through this kind of process corresponds nicely with the way that the spices are used.

From Box to Shelf

In England Robert discovers Indian cooking. For obvious reasons to do with the country's colonial past, Indian cook books, restaurants and spices were all readily available. More spices were brought and Robert's collection started to spill over into various drawers in his kitchen.

In 1982 the family buy a house and Robert purchases a small shelf for the new kitchen. The jars in the box were transferred to the shelf and placed in the thematic categories that had crystallised over the years (see step 4 of figure 2). Although this particular shelf is not the same shelf as the one described earlier in the paper (that one still being many years, and many meals, away) it can be assumed that it supported Robert's cooking in ways similar to the present day shelf.

In the transition from box to shelf there is a noteworthy qualitative shift that occurs. Whereas the chief function of the categories in the box was to aid Robert in locating specific spice jars (remember, only the top of the lids could be seen when the jars were in the box), the visibility afforded the jars when placed on the shelf gave rise to new, unplanned for, and unanticipated functionality, in addition to a general improvement of the previous functionality.[17]

Let me briefly outline some of the consequences of combining the categories of the box with the structural features of the shelf (the main differences between using the box and the shelf are also summarised in table 1).

Remembering which spices to use A cursory scan of the shelf can trigger memory of the spices included in a particular dish. When taking or replacing a jar on the shelf during cooking, adjacent jars may serve as reminders of spices still to be applied. In contrast, the content and labels of the jars kept in the box are not visible.
Finding a sought for spice jar A guessed at location of a particular spice is easier to confirm using the shelf, since feedback is instantaneous and category boundaries more distinct. Erroneous guesses are more costly using the box.
Correct replacement of spice jars Gaps are easier to spot and fill on the shelf and neighbouring jars also help to establish correct replacement. In the box there is a greater risk that jars will shift about, breaking up thematic groups.

Table 1. Summary of main cognitive differences between using the box and using the shelf.

Finding a sought for spice on the shelf has some similarity to finding it in the box: in both cases Robert is required to know the relevant category to which the spice belongs as well as the rough whereabouts of that category. One of the things that differentiates the two cases, however, is the ease with which the supposed loca-

17. Those with a fondness for evolutionary metaphors of artefact development (see e.g. Basalla, 1988; Ziman, 2000) might like to think of this event as a case of exaptation (Gould & Vrba, 1982). That is, as a feature which currently enhances fitness (i.e. cognitive congeniality), but which was not originally built for the role it now plays.

tion of the spice is then confirmed. On the shelf, feedback is almost instantaneous, the jars are stored one jar deep and can be easily scanned (compare this with the shelf in his first kitchen which was just a single plane). In the box, a jar has to be lifted before its identity can be confirmed. Not only is it easier to locate a particular jar on the shelf, but the cost of a faulty guess is much less compared with the extra effort incurred when picking out the wrong spice from the box. The categories in the box also have looser boundaries and are harder to pinpoint than the more rigid categories on the shelf. From just looking at the box the various thematic groups are not readily apparent and there is the constant risk that the location of particular jars will shift about during prolonged use of the box.

The removal of a jar from either box or shelf leaves a gap that can later aid in the replacement of the jar. On the shelf the gap is easy to spot and fill, and the visible identity of adjacent jars can confirm a correct replacement. In addition, adjacent spices can serve as reminders of spices still to be applied. In the box the identity of neighbouring jars can only be established by lifting them up above the level of the box.

Earlier in the paper it was noted that the vertical positioning of the shelf (in relation to Robert's body) also contributes to the ease with which particular spice jars are found and retrieved.

Reaping the Benefits

The ways in which the shelf supports cognition is a, mostly, unanticipated result of the combination of the categorisation of the spices contained in the box with the structural properties of the shelf. The claim made here is that the improved cognitive congeniality of the shelf is partly accidental, and historically contingent. However, the new ways of working afforded by the shelf are not automatically achieved. Although there are cases were a change in the material means of an activity entail a concomitant change in procedure, in this instance some adjustment had to be made before the benefits of the new set-up could be reaped. There may be many ways in which the shelf could potentially be used, but only some of these are an improvement over the previous use of the box. It is by using the shelf in particular kinds of ways that it is able to scaffold cognition. An example given earlier in the paper is a good illustration of this. You might recall how Robert's strategy (or habit) of taking down and replacing spice jars one at a time permitted the shelf to simplify the cognitive demands of the task. If several jars were instead taken down together, then Robert would be faced with the additional chore of pairing each of the jars with the appropriate gap.

The transition from spices kept in a box to having them arrayed on a shelf, the adjustments to, and appropriation of the resulting structure by Robert, is an interesting case in which an artefact (or artefactual complex) grows and develops in cognitive congenial ways. Needless to say, all artefactual change does not lead to improved functionality or to cognitive congeniality. Nevertheless, the process described may have greater generality than this single case.

Similar mechanisms can be found in, for instance, the gradual co-evolution of the form of books and bookcases (Petroski, 1999). Before the advent of the printing press books were rare and expensive luxuries, either locked up or chained to their bookcases. As they became more numerous, vertical partitions were introduced to the then standard bookcase design in order to prevent the shelves from sagging. Although the motivation for these partitions was originally structural the partitions later came to play an important role in locating books. Catalogues, usually posted on the end of a bookcase, grouped the books in accordance with the partitions that contained them. Even as late as 1749 catalogues were still not alphabetical, but based on these tables of contents.[18]

Concluding the Story

In 1988 Robert and his family moved back to Sweden. During the summer of their return they lived in a caravan and a selection of the spices were again back in a box. Later, when Robert and his wife moved into a flat, Robert put in an order for a new shelf. In 1990 they finally bought a house and the two IKEA shelves, described at the beginning of the paper, were purchased.

18. For some other interesting types of interactions and exchanges between co-located artefacts see de Wit et al. (2002).

At present almost all spices have been transferred to recycled Coleman's mustard jars. One could argue, from a cognitive standpoint, that a mix of jars, of varying appearance, would have been better (providing redundant cues as to identity), but here Robert prefers to let æsthetic concerns take precedence. Cognitive congeniality is, after all, but *one* factor that governs the shaping of our environment.

Today Robert has set a self-imposed limit on the continued growth of the shelf. He confesses to having been "a bit of a collector" in the past, buying spices in order to learn about them. Now he knows more about his needs and there is also a greater pressure on available space with Japanese and Thai cooking having been recently added to his repertoire.

Since the time of the original study, and the last and most recent interview with Robert, the shelf has undergone further change and is still in flux. Since the initial study was conducted Robert's wife has converted to using organically grown spices. As a consequence the bottommost shelf has been cleared for that use (as seen in figure 1). The two shelves above it now house Western spices used by Robert alone. However, a short while after Robert's wife converted to organically grown spices many shops in Sweden ceased stocking them and they have become increasingly difficult to buy. Robert's wife confesses to now using the ordinary spices on the shelves above hers to "top up" her own jars.

Will the organisation of the bottommost shelf persist, as a vestige of an ephemeral fad, or will the organisation of the shelf eventually return to the one described in the text? One thing is certain, the present shape and organisation of the shelf is unlikely to end here and will undoubtedly continue to change, in concert with Robert's unabated interest in cooking and in response to ever changing circumstances.

Discussion

The story told here is a reconstruction of events taking place during a period of roughly thirty years. As we have seen, the evolution of the shelf is intimately tied to the changing circumstances of Robert and his wife, Robert's intensified interest and growing knowledge of cooking, and even to changes at a societal scale in eating habits. The biography of the shelf and its use has allowed us to glimpse some interesting things, a few of which I think bear repeating.

One important insight is the realisation that much of the structure that supports cognitive activity may have partially *non-cognitive origins*. At least some structure seems to be the result of chance, circumstance, compromise, surrounding agents, and the shaping force of repetition. Perhaps the most significant moment in the present case study is the emergence of new and unanticipated functionality from the combination of previously unrelated structures.

I believe that insufficient attention has been paid to these sorts of processes and their impact on tasks and task environments. We need to continue to explore and expand our catalogue of these phenomena, but more pressing, perhaps, is the work of disentangling and understanding their interplay with other, more purposeful and intentional, processes. Some changes to an artefact or task environment can impact the cognitive ease with which a task is performed without any major changes in procedures or techniques, whereas other changes are accompanied by a concomitant transformation in the way a task is carried out. How people adapt to new structure, and appropriate and incorporate it into more congenial forms of a task, is a key part that needs to be properly worked out.

A better understanding of these phenomena may also serve as an important corrective to theories of design, production and artefact functionality, that are excessively intentional.

Another important and related point is the significance of use. One thing that has been demonstrated in this paper is that the cognitive congeniality of an environment is as much a function of an agent's particular use of that environment as it is a function of the environment itself. It is the particular ways in which things are used that *permits* them to contribute in cognitively beneficial ways. Cognitive congeniality is a relational property, and cognitive biographies must include both the changing forms and shifting uses of things.

As was noted in the introduction, the reconstructive nature of cognitive biographies may be the cause of some concern. The way that the present biography was created, for instance, was through repeated interviews coupled with study of the contemporary form of the artefact and associated activities. The biography is, by necessity, constructed after the fact; consequently there are aspects of it that are based on conjecture. Though cognitive biographies have some distinct advantages they also open the doors to speculation and presupposition.

This problem, however, is not unique to the present endeavour. There are a number of related enterprises that have learned to deal with similar issues. There is much that we may learn, for instance, from the history of technology (for reviews see Staudenmaier, 1984, 1990), the anthropology of technology (e.g. Lemonnier, 1986, 1992; Pfaffenberger, 1988, 1992), social construction of technology studies and actor-network theory (for a review of both see Bijker, 1994), as well as the diverse and numerous branches of archaeology. None of these disciplines (and there are more than those listed here) are specifically focused on cognition, but all are concerned with the processes behind changes in material culture. These areas may provide us with supporting evidence, complementary perspectives and methodological innovations and insights.

Recent focus on cultural biographies of objects (e.g. Appadurai, 1986; Kopytoff, 1986; Gosden & Marshall, 1999) is an interesting example, not only for the partial neologistic parallelism, but because of shared methodological issues. Both kinds of biographies seek to retrace sequences of relations between people and things. In the case of cultural biographies of objects it is a sequence of shifting meanings that is the elusive quarry, in the present case it is a sequence of uses and cognitions that needs to be reconstructed. Each quarry is as intangible and ephemeral as the other and we might find that there are methodological solutions to be shared.

Although interpretative science is difficult there are some potential rewards to be had. Cognitive biographies allow us, for instance, to explore real life events and changes spanning long time periods. And they allow us to concentrate our efforts on sites were significant change has already been established. But there is a further, fundamental reason for constructing cognitive biographies of things.

Tracing the history of a thing and its use can help us understand the present use of that thing. A cognitive biography allows us to better discern the cognitive roles currently being played by an artefact. Against the backdrop of earlier incarnations of an activity, and previous forms of an artefact, the cognitive functions of a thing are able to stand out in relief. For example: a feature of an artefact may be the result of a response or adjustment to problems inherent in previous versions of the task. Knowing about these earlier phases, enables us to discern (or, at least, to explore) the roles being played by this feature. Overlaying succeeding phases of an activity with preceding ones can often point us to possible areas of cognitive significance.

A disregard for the developmental trajectories of environments, tasks and people, will therefore lack some of the essential ingredients necessary for a genuine understanding of the cognition of task performance.

Acknowledgements

I would like to thank the participants and organisers of the Doing Things with Things workshops in Copenhagen for their reactions to the ideas that eventually became this paper, and the participants of the Lund University Cognitive Science seminar for their comments on an earlier draft. I am also grateful to Björn Nilsson for numerous suggestions, and Pieter Vermaas for some particularly cogent points of criticism. Thanks are also due my informant Robert for his help, as well as for his delicious cooking. This work was supported by the Swedish Foundation for Strategic Research.

PAPER FOUR

BUILDING THOUGHT INTO THINGS

ABSTRACT: **This paper retraces the steps of Bærentsen's (1989) cognitive analysis of the evolution of the rifle. Bærentsen's central thesis is that the actions and thought processes required to operate a rifle, at any one stage of its development, are "built into" subsequent generations of the artefact. In the process of retracing and critically evaluating Bærentsen's analysis, a slightly different view gradually emerges, in which greater attention is paid to the interplay between the physical properties of artefacts and the structure of tasks. The interdependencies between artefact design, task structure, task goals and cognitive tasks are explored and developed throughout the paper.**

INTRODUCTION

Can thought be built into things? The question sounds like something a child might ask you, not because there is something inane about the query, but because we are asked to consider the merger of categories we have learned to keep separate. A first attempt at wringing some sense out of the question stumbles over the almost paradoxical coupling of mental and material realms. It is hard to imagine what it would mean for thought – perception, reasoning, judgment etc. – to be combined with, or built into, material objects. This tension can be felt even without recourse to the philosophical mind–body problem.

It is perhaps the apparent incongruity of the idea which lies behind the allure of the many thinking objects we find in modern mythology, be they mirrors with opinions on the things they reflect, or spaceships with dysfunctional personalities.

The second of these fanciful examples corresponds to one avenue in which the question has already been extensively explored: the field of artificial intelligence, or AI. This will not be the subject of this paper. Instead we will search for traces of thought in *everyday objects,* artefacts that were not created with the express intent of duplicating thought.

To this purpose I will retrace an analysis carried out by Bærentsen (1989) in which he examines the gradual objectification of thought in the evolution of the rifle. Bærentsen's central thesis is that the actions and thought processes required to operate a rifle, at any one stage of its development, are "built into" subsequent generations of the weapon. I propose that we call the kind of intelligence we are looking for in everyday objects "artefactual intelligence," to distinguish it from the kind of intelligence studied in the field of AI.

The historical development of firearms is an ideal material for our purpose in many ways. First of all, Bærentsen has done a fine job of summarising the developmental history of the rifle, and his thesis is both original and consistently pursued. Secondly, the artefacts and tasks examined can be understood without any specialised knowledge, and are sufficiently circumscribed to permit treatment in the present format.

In the process of retracing and critically evaluating Bærentsen's analysis, an alternative view gradually emerges in which greater emphasis is placed on the interplay between the physical properties of artefacts and the fine-detailed structure of the tasks performed.

The evolution of the rifle

The first documented hand-held firearms appear in the mid fourteenth century. These consisted of little more than simple metal cylinders attached to wooden shafts, with the far end of the contraption supported by a stand. The guns were loaded with gunpowder and used to fire a bullet or a handful of stones. Discharging the gun was achieved by plunging a hot iron rod through a small hole in the top of the cylinder and into the gunpowder.

Artefact design and task structure

In his analysis, Bærentsen draws attention to two crucial aspects of the task of firing the earliest guns that are imposed on the task by the design of the artefact. Using a hot iron rod as the means of igniting the gunpowder necessitated the close proximity of the gunman to a fire or a brazier in which to periodically heat the rod. The use of the hot iron thus greatly curtailed the movements of the gunman, and consequently, the positioning of the firearm. The second aspect of the task, imposed by the design of the gun, pertains to the actual operation of the device.

Putting the hot iron through the hole in the gun placed great attentional demands on the person using the weapon: attention being locked into finding the hole and into monitoring the motor actions required for fitting the hot iron into it. A consequence of this is that little attention is available for other, simultaneous aspects of the task, most importantly, the aiming of the gun at a chosen target. Due to attentional demands posed by the properties of the gun, aiming and firing become two incompatible actions which had to be carried out in succession. This is in contrast with contemporary firearms in which the two actions can be carried out simultaneously. Although the immediately following generation of firearms did not solve these attentional problems, it did increase the mobility of the gunman, as we shall see.

Building practice into things

The next design innovation, appearing sometime during the fifteenth century, is the replacement of the hot iron – as the means of setting off the gunpowder – with a slow-burning fuse. With the fuse (also know as a match), the gunman was no longer yoked to a source of heat, but able to carry the gun with him over greater distances as well as on horseback. The fuse, and the resulting portability of the gun, entailed further design changes in the gun itself, but it is the fuse which is predominantly responsible for the radical transformation of the structure of the task.

Bærentsen suggests that the brazier, which was used for heating the iron rod, as well as the operations required for servicing it, were in some sense "built into" the fuse. Although the claim is suggestive, it is not exactly clear in what sense it is to be understood. I will not try to extricate the meaning of the phrase as it occurs here, as I believe that in this instance it acts to obscure a slight misanalysis of the changes brought to the artefact and the task.

Changing the timing of the task

It seems unreasonable to suppose that the brazier has been "built into" the fuse, however metaphorical we choose to be, since surely even the fuse has to be lit at some point and at some source. In a way similar to the original hot iron, it too is an intermediary between a source of fire and the gunpowder which it is used to ignite. It cannot solely be the portability of the fuse that enables the greater movements of the gunman, for the hot iron is also, at least in principle, portable. The iron would perhaps be unwieldy to carry on horseback, and also associated with a certain risk – such as the involuntary branding of both rider and horse – but the real problem is that the iron would cool off all too quickly, rendering it useless for the task at hand.

Here is the significant difference between the two artefactual complexes, accounting for the changes in the way that the task is performed. It is the fact that the fuse is able to retain a particular kind of state (i. e. combustion) necessary for subsequent operations, and that it is able to retain this state for a sufficient amount of time, whereas the hot iron soon turns into a cool iron, that accounts for the acquired mobility of the gunman. The slow combustion of the fuse allows the gunman greater flexibility in choosing *when* to fire the gun. It is this change in the timing of the task, in turn, which gives the gunman the greater freedom of movement.

The ability of the slow-burning fuse to maintain a prolonged state of combustion might be thought of more abstractly as the ability to retain a particular state necessary for subsequent operations of a task. Couched in these terms this is clearly a property common to artefacts other than those analysed here, though varying from case to case in both longevity and dynamics. If there is a general account of the effects of this property on the timing and sequencing of tasks, it is one that will have to be provided elsewhere. In the interim, let me note that although the capacity of an artefact to retain a particular state may be associated with flexibility in the timing and sequencing of a task, as we have just seen, there are tasks in which a transitory state of the artefact is absolutely essential. Using a bow and arrow is, if perhaps not the ideal example, one such task: the successful operation of the bow depending entirely on its elasticity.

It has not been my intention to admonish Bærentsen, the transition covered – the replacement of iron and brazier with the slow fuse – take up no more than a few lines of the original paper. However, rehearsing this part of Bærentsen's analysis has allowed us to see the profound effect on task structure that may result from the substitution of just one artefact, or more correctly, one component of an artefactual complex. We have also seen how the nature of the change in the task was closely coupled to a particular property of the substitute artefact. This general view of the relation between the properties of artefacts and the tasks performed with them is one we derive from Norman (1991).

Building-in sensorimotor operations

We have seen how the fuse helped to solve the problem of the gunman's restricted movements, the next design innovation was one that served to alleviate some of the difficulties inherent in aiming and firing the gun. You will remember that aiming and firing the first hand-held firearms were actions that were largely incompatible due to limits on attentional resources and that as a consequence aiming was mostly left to chance.

In the following generation of guns the fuse was mounted on the end of a mechanical arm running along the length of the barrel (see figure 1). Viewed from the side the mechanical arm was approximately z-shaped, with the top and bottom strokes of the letter being greatly elongated, and the whole structure pivoted at the centre of the diagonal stroke. Pulling the near end of the arm towards the barrel caused the far end, to which the burning fuse was fastened, to tip towards the hole, igniting the gunpowder. Firing the gun no longer required the kind of attentional resources previously needed as the arm could in principle be operated without looking. Attention could now instead be devoted to the different components of the task of aiming.

Figure 1. Firearm equipped with a mechanical arm.

What has occurred here, Bærentsen tells us, is the objectification of the sensorimotor operations that previously made up the action of discharging the firearm. According to Bærentsen there are two aspects of the performance of the task which can both be said to have been "built into" the material structure of the gun. One aspect is the sequence of motor actions itself – the physical movements of bringing the fuse to the aperture (the touch-hole) on the top of the gun. Bærentsen argues that these have been taken over and are now performed by the mechanical arm. The second aspect pertains to the perceptual and cognitive operations involved in finding the touch-hole, and in guiding and monitoring the motor actions of the task. The claim here is that these cognitive processes have been "fixed," once and for all, in the physical structure of the artefact.

How should we understand these claims? What might it mean to "build in" or to "materialise" motor operations, or to "fix" perceptual and cognitive operations? Lets begin with the first idea, the idea that the actual movements of the gunman are somehow captured by, or solidified in, the material structure of the mechanical arm.

BUILDING MOTOR OPERATIONS INTO THINGS

There is a sense in which the mechanical arm does seem to have taken over the physical actions that were once performed by the operator of the gun. If we focus only on the movements made by the fuse, and blot both gunman and mechanism from view, at least part of the trajectory of the fuse will be approximately the same, regardless of whether the fuse is mounted on a mechanical arm or gripped in a human hand.

At first sight it may indeed look as if the mechanical arm is mimicking the movements of the gunman's hand. However, we must not forget the efforts of the gunman who we momentarily removed from view. Without the gunman to pull on the arm, supplying the necessary energy and momentum, it would simply remain inert. Bærentsen is quite clearly not suggesting that the mechanical arm supplants the actions of the gunman, it must be just as obvious to him that the agent has to initiate the action, as well as provide the necessary muscular force. The separation of motor operations from the cognitive control of those operations is surely a distinction made for analytical convenience. But with the source of power extraneous to the arm, and the issue of control separated from that of execution, the question remains as to what it is that the arm has taken over: what it is that has been built in.

A set of five interrelated transformations

There may still be some way of salvaging Bærentsen's claim, some way of making sense of the idea of the mechanical arm doing the work of the flesh and blood arm, but the details available in the original analysis are too meagre to work with. At this point I think it would be more fruitful to present my own conceptualisation of what the introduction of the mechanical arm will have entailed. I propose that we view the adoption of the new gun design as a set of five interrelated transformations.

1) The Physical Structure of the Artefact Transformed

The change in the design of the gun is itself a kind of transformation: a transformation of the physical structure of the artefact. In some ways this is the most obvious and readily apparent of the transformations to be described. Bærentsen's reconstruction of the evolution of the firearm is largely conjecture based on just such a sequence of physical transformations.

2) The Structure of the Task Transformed

The physical changes that are undergone by the artefact are accompanied by a transformation of the structure of the task in which the artefact is employed. In the case at hand, instead of having to locate the touch-hole on the top of the barrel of the gun and guide the fuse into it, the gunman has now only to pull on the mechanical arm until the gun fires.

Whereas these first two transformations are changes that take place from one generation of the artefact to the next, the third and fourth transformations below are ones that operate on the occurrent actions within the task.

3) The Performance of the Task Transformed

The mechanical arm, in addition to presenting the gunman with a new kind of operation to perform, also serves to transform the actual performance of the operation. The physical constraints inherent in the artefact will both shape and guide the actions of the gunman whilst the arm is pulled on. The mechanical arm will resist certain movements – such as any attempt to pull the arm sidewise – and facilitate others – for instance, those movements that conform with the direction in which the arm is hinged, but also movements that are oblique to these. The agent operating the arm will not be impervious to the affordances offered, and will naturally respond and adjust to these, but the point here is that the arm will to some extent *forcibly* alter the direction of the manipulations of the person operating it.

4) The Gunman's Movements Transformed

As the movements of the gunman are transmitted by the arm, their direction and amplitude will be altered and transformed by the mechanical properties of the artefact. The upward movement of the gunman's hand, for instance, will be transformed into a downward movement of the fuse. The extent of the transformation of the amplitude of the hand movement will depend on the placement of the pivot and the length of the various parts of the arm.

This fourth transformation is a transformation of the vector provided by the gunman. The gunman's input is itself, of course, a product of the interaction between the agent and the artefact, as was detailed in the transformation described under the previous heading.

We can now note some ways in which the transformation of the structure of the task by the artefact is dependent on the artefact's ability to shape and transform the actions *within* the task. One thing the mechanical arm does is to greatly simplify the task facing the gunman. The arm has only one degree of freedom and this makes the attendant operation very simple. More important, almost all variation in the gunman's input will be absorbed by the arm, and channelled into an appropriate movement of the fuse at the other end. The arm acts much like a funnel – both in a metaphorical and in an analogical sense – channelling a wide range of actions into one efficacious movement. Since the task is also fairly insensitive to variations in timing, almost any tugging or pulling on the arm will be translated into an appropriate effect. The skill needed to operate the arm is therefore very slight.

The real elegance of the design is not that it presents the gunman with a simple task to perform, this would be of little value if the demands of the task were not at the same time met, but that it enables a simple manipulation to satisfy task requirements. It is because the mechanical arm is able to translate one kind of movement into another, that it is also able to change the structure of the task of firing the gun.

In other words, the mechanical arm presents a simple task that requires little skill to perform, and the simplicity of this task is dependent on the properties of the mechanical arm. In addition, the properties of the arm ensure that the simple action required of the gunman will have the requisite effect.

5) The Goal of the Task Transformed

Finally, these changes to the task and artefact will, I hazard to guess, have been accompanied by a concomitant transformation of the goal of the task. With attention released from the sub-task of discharging the gun it could now be employed in other activities such as improved aiming. When using the previous generation of guns it would have been expedient to aim at clusters of people in order to increase the chances of a hit: the choice of a larger target partly compensating for limitations in precision. With more attention available for the sub-task of aiming, the choice of target could be made more specific. Instead of hoping to just hit *someone,* the gunman could now select a specific target and aim to hit the chosen target. The introduction of the mechanical arm did not simply improve the gun, it also changed the kinds of task for which it was used. At a very general level of description the task has, of course, remained the same – wounding the enemy – but at a more specified level the task has become quite different.[1]

Being able to target a specific individual not only changes the structure of the task of firing a gun – the act of aiming and firing – but also influences the strategic uses to which the gun is put. This relatively small change in the artefact may thus have a large impact on numerous other tasks, as well as on the relationships between these tasks.

---

1. The descriptions of the task are those of an outside observer and there is always some indeterminacy as to how the task is conceptualised by the agent himself (Draper, 1993).

The coevolution of tasks and artefacts has been called the task–artefact cycle (Carroll, Kellogg & Rosson, 1991). Tasks set requirements on the design of suitable artefacts. The resulting artefacts, in turn, suggest novel possibilities and impose new constraints, which taken together, often redefine the task for which the artefact was originally developed. The new task then sets new requirements and the cycle continues.

Fixing perceptual and cognitive processes

We should now be in a better position to judge Bærentsen's claim that certain cognitive processes have been "fixed" in the physical structure of the mechanical arm. The two things which have supposedly been fixed in the artefact are: the perceptual process of locating the hole on the top of the gun, and the cognitive process of guiding the fuse into the hole.

I am inclined to agree that there is a sense in which the process of finding the touch-hole has been set or fixed in the structure of the device. Whoever manufactured or calibrated the mechanical arm would have had to locate the hole in the top of the barrel – a process involving the perceptual mechanisms of the artisan responsible – and then set the arm appropriately. Thereafter, the mechanical arm will guarantee that the hole is "found" every time the arm is pulled. To be able to operate the arm, it too must be found, of course, but its size and placement makes this an automatic consequence of handling the gun, and not chiefly a perceptual process.

Our intuition in this case seems to rest on the unspoken assumption that to qualify as having being "fixed," a cognitive process has to be appropriately exercised during manufacture. There is perhaps a notion of "fixing" for which this would be a necessary criteria. By this criterion however, the cognitive process of guiding the fuse into the touch-hole will not qualify as having been fixed in the structure of the artefact. Whereas setting up the mechanical arm so that it will find the hole in the gun will require that the artisan first finds the hole, setting up the arm so that it will guide the fuse appropriately does not require the artisan to exercise the processes corresponding to the one the arm will be replacing. Making and placing the arm will involve numerous perceptual and cognitive judgments, but not the rehearsal of those performed by a gunman manipulating an unattached fuse.

Moving away from the *genetic* criteria above we might instead note the particular way in which the mechanical arm replaces the cognitive processes of guiding and monitoring the movements of the unattached fuse. There is a part of the task, a certain task requirement, which remains constant through both generations of the gun. In both generations, the fuse has to pass through a certain trajectory. Or put differently, the fuse has to be brought from a state in which it is not in contact with the gunpowder to one in which it is. In one case the movements of the fuse are controlled by mainly cognitive processes, in the second case, most of the control is achieved through the physical constraints inherent in the artefact. Despite these differences there is a sense in which the two are performing the same job.

This is a different kind of conception than the genetic interpretation of what it might mean to "fix" cognition or perception in material things. Here we focused mainly on the nature of the task that was being solved. There is likely to be other interpretations of the notion. We could concern ourselves, for instance, with the mechanisms involved in performing the task. For the time being I will content myself with having pointed out a couple of possible interpretations.

Two ways of dividing one's attention

The next change brought to the design of hand-held firearms is the introduction of the first spring-loaded locks. Instead of pulling on a large and cumbersome mechanical arm, the gun is now discharged by means of a delicate lever which releases a small spring-loaded arm (see figure 2). Bærentsen's take on this is that the last remaining vestiges of attention, that were demanded by the mechanical arm, have now been released. Complete attention could finally be devoted to aiming the gun. As a result the stock of the gun changed shape, making it possible to rest it against the shoulder, thus further transforming the task of aiming. The introduction of lock and stock will also have made it easier to hold the weapon steady at the moment it was discharged, which would have greatly contributed to the weapon's accuracy.

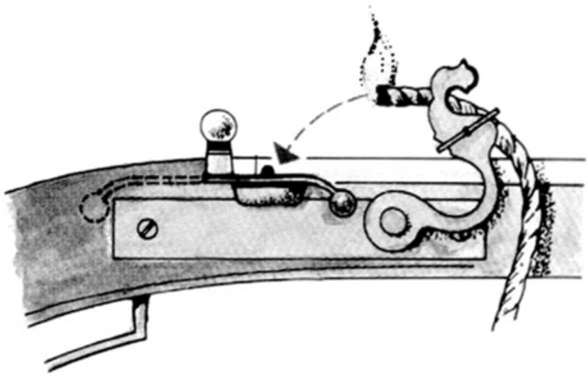

Figure 2. Matchlock mechanism.

We must be careful lest we think that the attentional demands placed on the agent have simply been blown away by the introduction of the spring-loaded lock. Bærentsen is of course correct that attentional demands have been lessened, it so happens, however, that they turn up in a different place.

Like the mechanical arm, the spring-loaded lock is still dependent on the gunman for its motive power, but now in a more indirect and less immediate fashion. Motive power is supplied by the gunman when the gun is cocked, and stored in the mainspring until the time at which the trigger is pulled. Priming the gun requires the attentional resources of the gunman, but at a time when these are not being competed for. The ability of the spring to store energy (another instance of the ability to retain a change of state) solves the problem of limited attentional resources by redistributing the demand for these resources. The artefact permits part of the action of discharging the gun to be performed ahead of time, *storing* the action if you will. The action can then be deployed at a time when the gunman would not normally be able to perform the action in a satisfactory manner.

This is an interesting general strategy for handling limited attentional resources and can be compared with the slightly different strategy which lay behind the adoption of the mechanical arm. Whereas the spring-loaded lock is able to redistribute the demand for attentional resources over time, the mechanical arm of the prior generation of guns solved a similar problem by redistributing an attentionally demanding tasks over modalities.[2] With the first generation of guns, both aiming and firing made claims on visual attention. The mechanical arm was able to shift most of the attentional demand of firing the gun to a neighbouring channel.

Rationalising task structure

Successive innovations in the design of the triggering mechanism – the wheel-lock, the snaphance lock, the flintlock, and the percussion lock – are all able to produce a triggering spark or explosion on demand, obviating the need for a burning fuse. This further simplified the use of the weapon, an entire step – that of lighting the fuse – being removed from the process of operating the gun.

According to Bærentsen the complex sensorimotor operations that were previously involved in servicing the fuse have partly been made redundant, and partly materialised, in the physical structure of the triggering mechanism. Again, I think that we can see this as a case of the artefact transforming task organisation.

The kind of transformation that I believe has occurred here is a form of rationalisation. Two actions which were previously separate – bringing the fuse in contact with a source of fire, and bringing the fuse in contact with the powder charge – have been brought together in a single action. Elements of the first action that are duplicated by the second action have been excised. That this is possible is due to the unique properties of the triggering mechanism – specifically the properties of the flint or pyrites – to convert a motive force into a spark.[3]

---

2. The spring-loaded lock actually combines both of these strategies.
3. I am simplifying the description and analysis of the process in order to bring the general idea across. In actual fact the locks which came to replace the matchlock did not ignite the powder charge directly, but via a small quantity of priming powder which first had to be deposited in the pan of the lock.

As with the previous transitions, the transition from mechanical arm to spring-loaded lock will have entailed changes to both structure and goal of the task. It will also have entailed changes to the structure and goals of related or even, what appears to be, unrelated tasks. The repercussions are vast and complex, and almost impossible to track. One interesting effect of the gradual move away from the use of a burning fuse is that the gunman was no longer as conspicuous during night or twilight hours. The momentary flare of later lock mechanisms being harder to track than a glowing row of smouldering fuses. This will have had clear tactical consequences.

Reorganising cognitive task structure

An aspect of the task of operating the earliest guns which has barely been touched on, but which is by far the most time consuming and attentionally demanding of the processes involved in servicing the gun, is the loading and reloading of the weapon. This was probably the most important single factor dictating the possible tactical uses of the gun. The exact method used to load the earliest firearms varies from design to design, but generally involved the deposit of a measured quantity of gunpowder and a shot down the muzzle of the barrel. After the shot had been fired, and before a second shot could be discharged, the barrel of the gun had to be cleaned out and the loading repeated.

It is interesting to note a transformation of the task which is not a direct result of a change brought to the design of the artefact. Bærentsen reminds us of the importance of the soldiers' training. One thing soldiers are made to do is to practice ritualised sets of movements, including those of loading, aiming and firing guns. This program of training transforms the task of loading the gun in at least two ways. First of all, the movements learned have been carefully pared down to essentials. This is one way of reorganising the structure of a task without transforming the physical artefact. Secondly, the repeated drilling of these sets of gestures eventually makes them automatic and almost effortless to perform. We could see automatisation as a kind of mental reorganisation of the task (cf. Hutchins, 1986). Outwardly, the gunman might seem to be performing the same task, say loading a firearm, but the inner organisation of the cognitive processes employed may have changed.

COGNITIVE HOMEWORK

Fine tuning and automatisation of the execution of the task minimises the time the gunman spends loading the gun and clearing out the bore of the barrel. But, regardless of how much time is spent in training, there is still a limit to how quickly the gunman can perform the task. There is one way however, in which the gunman's time *off the battlefield* can be used to speed up the loading of the firearm.

The procedure of loading the gun is made up of a number of separate activities. The gunman has to deposit a measured quantity of gun powder down the muzzle of the firearm, followed by a shot wrapped in a small patch of cloth. The powder, shot and patch all have to be brought out from various containers and coordinated in the task of loading. We should not forget to mention the fact that this often had to be done in the heat of battle.

Eventually someone had the brilliant insight that part of this process could be carried out in advance. The sixteenth century therefore sees the introduction of pre-packaged cartridges containing a shot and measured quantity of powder. These paper packets could be placed in the breech of a gun with altogether far less ado. The packages could be prepared by the gunman at his leisure, well ahead of the time when they were to be used. Being able to divide up the performance of the task not only made it possible to distribute the task across time, but also across people (cf. Hutchins, 1995a, on the notion of precomputation).

For Bærentsen the paper cartridges are yet another case of "fixing." This time rather a lot of things are being packed into the paper cartridge along with the shot and powder. As well as the motor operations of loading the gun we find the deliberations, decisions and perceptual processes which used to form part of the activity.

There is rather a lot to unravel here. I will not however, fully reanalyse what is going on. Loading the paper cartridges still requires certain motor actions and cognitive processes. We do not get around these by wrapping them up in a paper package. I think that the main thing to focus on is how the task has been restructured in time.

PRECOMPUTATION OR FOSSILISED PRACTICE?

Following these first few steps in the development of the rifle are countless innovations to every possible aspect of the gun's design. Improvements to the mechanism of the gun further simplified the process of loading and clearing and also increased the range of the weapon. Loading and clearing the gun became increasingly mechanised over time. Means were eventually found of exploiting the energy of the gun's recoil and using this for the automatisation of loading and clearing the breech.

The range of the gun, and how fast it can be reloaded, is of little consequence if the target can not be hit because the gunman's aim is too poor. There were other changes to the design of the gun which improved the support of the task of aiming. The introduction of the shoulder-stock is a change that has been mentioned already. By resting the stock against the shoulder, the gunman could look down the length of the barrel – instead of seeing it from a slightly elevated angle – making it easier to judge the direction in which the gun was pointing. The change in the design and handling of the artefact transformed the cognitive task facing the agent: lining up the muzzle of the gun with the chosen target is a predominantly perceptual task, whereas the previous task required a more complex set of judgments (cf. Clark, 1997, on the generality of this type of transformation).

The perceptual task of aligning the gun with a target was eventually simplified by fixing a small projection near the muzzle of the gun. This innovation was later refined by the addition of a second projection at the breech end of the firearm. The second projection, nearest the eye of the gunman, was cleft so that the first projection could be seen through it when the two were properly aligned. Seeing the far projection through the cleft ensured that the gun was being aimed straight.

Though these developments to the sight of the gun transformed the task of aligning the gun with a target, a mental correction still had to be made for the vertical deviation of the shot caused by gravity. To compensate for the effects of gravity on the shot, the gunman had to aim slightly above the intended target. How much above, depended on the distance from the target and the velocity of the projectile. The continued development of the cartridge made the velocity of the projectile an invariant factor. With velocity held constant, estimating the vertical deviation of the shot was a matter of judging distance alone, and of being able to transform distance estimates into appropriate adjustments to the aim. I am assuming here that shooting practice engendered a *feel* for the relation between distance and displacement, and that the gunman did not consciously calculate the displacement, or consult a booklet of printed tables.

In a later development of the sight the displacement of the shot at various distances is *stored* in the structure of the artefact and does not need to be remembered or calculated by the gunman.[4] The cleft was made part of an adjustable scale on which it could be moved up and down (see figure 3). If the cleft is set high on the scale – i. e. pulled in an upwards direction, away from the top of the barrel – the gunman has to tilt the muzzle upwards in order for the two projections to align (see figure 4). The markings on the scale were spaced to ensure an operative angle of the gun for every distance marked on the scale.

We can imagine several different ways in which the scale could have been calibrated. One way would be to take the gun out – or one just like it – to the shooting range and calibrate the scale through a process of trial and error. Alternatively, the spacing of the markings could have been calculated with pen and paper. Whichever method was used will not be apparent from just looking at the weapon.

4. The description and workings of the sight are adapted from Bærentsen. For some reason however, he omits to draw attention to the practice (i. e. that of calibrating the scale) materialised in the structure of the sight.

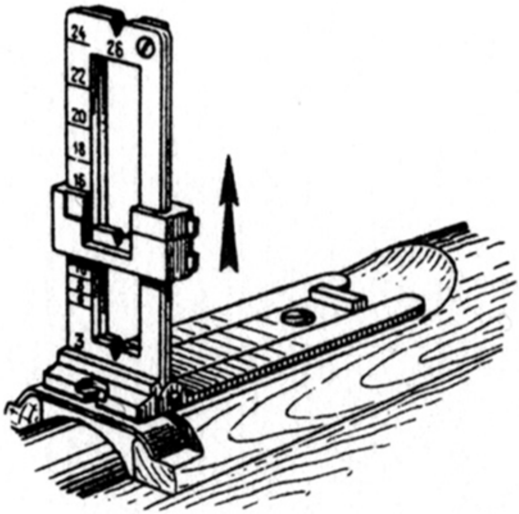

Figure 3. Adjustable sight.

So what exactly, if anything, has been built into the scale? Is it a repository of practical experience, or is it an external memory for a set of computations? This depends partly on the history of the particular device we are looking at. There is a sense in which both can be considered repositories of practice, although in more or less direct ways. The hypothesised “mathematical” means of calibrating the scale is not untarnished by practice, but relies on procedures and tools that are themselves the products of accumulated practice and experience.

Cultural choices

One of the things commending Bærentsen’s analysis is his sensitivity to the interplay between practice, cognition and artefact design. In this paper I have tried to examine and scrutinise this interplay in greater detail: moving a little closer to the physical properties of the artefacts under consideration, and looking a little closer at the fine structure of the tasks in which we find them. There is however one corrective with which I would like to end.

A backbone to this paper has been the historical development of the rifle. With Bærentsen's help we have traced the evolution of this category of artefact for a part of its history. The analysis has been more inclusive than is usual in discussions of technology, involving lines of influence that weave together interactions between artefact design, cognition, task structure and task goals. Although we have followed several causal chains it is important to point out that we have not been describing a determined system.

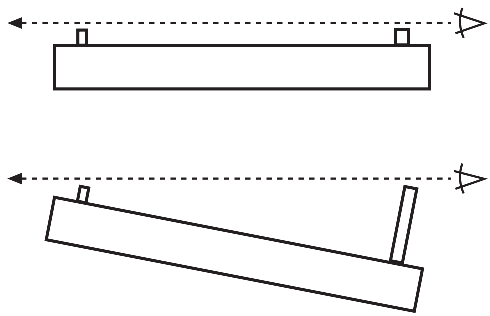

Figure 4. Schematic illustration of the adjustable sight compensating for bullet drop.

With hindsight it may seem as if the evolution of an artefact has a direction, that design improvements accumulate over time, slowly shaping the artefact into its definitive form. Is it not an incontestable fact that modern firearms are superior to the simple guns with which we have been preoccupied? Perhaps. It may be tempting to imagine the task–artefact cycle as a kind of spiral which gradually approaches an optimum solution, winding its way towards an ideal form of the task and artefact.

It would be hard to deny that there is an accumulation of innovations in almost any artefact we choose to look at. The latest generation of an artefact category is often the direct response to the perceived shortcomings of the immediately preceding generation (Petroski, 1992).

However, there are few tasks which ever reach a definitive and unalterable form. One reason is that surrounding activity rarely remains constant or unaffected by the evolution of the task. Take the rifle as an example. A plausible progression would seem to be towards deadlier and more accurate weapons – a scenario which is partly vindicated by the history of the modern rifle – but there are also limits to pure effectiveness. Weaponry that is too effective may be incompatible with other goals, for instance moral, political, economic or even strategic goals.

Another limit to pure effectiveness is that rifle design, like all design, is by necessity a trade-off between the effectiveness at different types of tasks that might be performed with the weapon. What is essential to one task is sometimes an impediment to another. This is part of the reason for the great diversity of rifle designs, with special rifles for specialised tasks. The occasional branching of the developmental line of an artefact, with new cycles spinning off in different directions, is one thing which the task–artefact cycle fails to include.

There are plenty of instances when effectiveness is shunned for personal or cultural reasons. In the Upper Verdon valley in the south-east of France for instance, the most experienced hunters are expected to hunt with a small-bored rifle – normally only used by women and youngsters – to demonstrate their superior proficiency (Govoroff, 1993). This is only one of many examples of what Lemonnier (1992) calls "technological choice": cases of people choosing particular technologies, or parts of technologies, for other than purely instrumental reasons. Cultural influence on technology is more pervasive than we might at first think. An important influence on how any given piece of technology develops is how that technology is conceptualised in the culture which supports it (cf. Quilici-Pacaud, 1993). The entrenched conception of what a gun is and does, for instance, is probably one of the most important determinants of the direction the development of the gun will take. There are design solutions which are perfectly valid on instrumental grounds, but which would never be pursued simply because the result would not look sufficiently gun-like.

ACKNOWLEDGEMENTS

My sincere thanks to Henrik Gedenryd for feeding me articles and ideas, and to Peter Gärdenfors for his patience and for his comments on the penultimate draft of this paper. Countless thanks go to Björn Nilsson for our numerous and stimulating discussions on topics touched on herein.

PAPER FIVE

THE FUTURE OF SELF-CONTROL: DISTRIBUTED MOTIVATION AND COMPUTER-MEDIATED EXTROSPECTION*

ABSTRACT: In this article we discuss the social implications of context-aware computing technology for the perennial human problem of self-control. We present a basic model of the domain of self-control, and provide a range of suggestions for how modern sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities. The model consists of two core concepts. The first we call Computer-Mediated Extrospection, which builds upon the familiar idea of self-observation or self-monitoring, and concerns itself with the crucial need for accumulation and explication of self-knowledge in any rational person-centred decision process. The second concept is Distributed Motivation, which we see as a natural extension of the idea of precommitment and self-binding that is often discussed in the self-control literature. The article ends with a discussion of issues of flexibility, and ethical concerns about privacy and persuasion in possible context-aware applications for self-control.

1. INTRODUCTION

The ubiquitous vision is one in which computers take an increasing part in our everyday activities, in ways that mesh naturally with how people think, act and communicate. We are excited by this vision, but feel that the full possibilities offered have yet to be explored. Work within ubiquitous computing and context awareness has made us increasingly familiar with computers that mediate our interactions with the world, *but what about computers that mediate our interactions with ourselves?* We believe that com-

*Written together with Lars Hall and Petter Johansson puters can be made more powerful by letting them gain information about the user, but in a similar manner we also believe users can be made smarter and more powerful by letting them gain additional knowledge about themselves. To this end we will here propose some ways in which sensor and computing technology might be used for purposes of *self-control*. This is an avenue that, surprisingly, has remained largely unexplored.

Recently, HCI researchers have shown a growing interest in the motivational role that computers might serve in human conduct – i.e. what has become known as *persuasive computing* (Fogg, 2003). However, this field is still in its infancy, and very few explicit connections to theoretical or empirical research on self-control have been made (e.g. to important work like Rachlin, 2000; Elster, 2000; or Ainslie, 2001).

The lack of qualified research on the possible use of information technology to alleviate problems of self-control is, we feel, a very unfortunate state of affairs. Issues of self-control are extremely pervasive in modern societies. Take for example the use of tobacco. According to the latest *World Health Report* (2001) our planet harbours more than 1.2 billion smokers worldwide and tobacco accounts for well over three million annual deaths. Costs are more difficult to calculate, but a recent World Bank report on the economics of tobacco control estimates that in high-income countries smoking-related health care account for 6–15% of all annual health care cost (World Bank, 1999). Similar figures can be found in relation to regulation of dietary intake. As measured by the standardised Body Mass Index we now have roughly 1.1 billion overweight people in the world (Gardner & Halweil, 2000). In the US alone an estimated 300 000 people die each year of causes related to obesity (Mokdad et al., 2001). However, problems of self-control and self-regulation are not only operative in such salient and life-threatening domains as craving and addiction, but also in the minute workings of everyday plans, choices and actions. Ameliorative action is as pertinent to the dreadful experience of withdrawal from heroine, as it is to innocuously hitting the snooze-button on the alarm clock, and missing the first morning bus to school (Rachlin, 2000; Ainslie, 2001).

The purpose of this article is to present a succinct model of the domain of self-control that maps both the degree of severity and the great variety of self-control problems, as well as possible remedial actions using modern sensor and computing technology. The model consists of two core concepts, or tools, that we believe may serve an important role in elucidating the problem of self-control from a ubiquitous computing perspective. First, we introduce the concept of *Computer-Mediated Extrospection*, which builds on and expands the familiar idea of *self-observation* or *self-monitoring*, and concerns itself with the crucial need for accumulation and explication of self-knowledge in any rational person-centred decision process. Secondly, we present the idea of *Distributed Motivation*, as a natural extension of previous discussions of *precommitment* and *self-binding* in the self-control literature.

Issues of context awareness occupy centre stage in the field of ubiquitous computing and human–computer interaction (Dey, Abowd & Salber, 2001). The most relevant aspects of context are also generally agreed upon. For instance, according to Dey, Abowd and Salber (2001): "Context is typically the location, identity, and state of people, groups and computational and physical objects," where state refers to "physical, social, emotional, or informational state". In practice however, modelling context has been mostly confined to information about identity and location, and far less attention has been paid to the psychological states of people. For purposes of research and development this has been a pragmatically sound strategy, and more recently it has also been amply demonstrated how pertinent contextual constructs can be leveraged from the combination of simple environmental measures (see e.g. Gellersen, Schmidt & Beigl, 2002). However, if the computer is to become a tool for augmenting interactions with ourselves these kinds of measures must be combined with more intimate, psychophysiological measures. In the kinds of applications we discuss in the present paper the affective and cognitive states of users, and the emotional context of interaction, are essential and inescapable aspects of context.

In presenting our model, we draw upon existing research in ubiquitous computing and context awareness (and from conceptual neighbours like wearable computing, telemedicine, affective computing, and the aforementioned field of persuasive computing), but to make our points clearly we also include references to future scenarios and hypothetical cases. It is our hope that the model and our discussion will provide a principled and useful way for designers of human–computer interfaces and context-aware systems to approach the domain of self-control, as well as to provoke further debate on the possible role of computing technology in matters of human motivation.

The outline of the article is as follows. First we present an overview of the problem of self-control, and then, in consecutive sections, we introduce and discuss our two conceptual tools. The article ends with a discussion of different self-control scenarios drawn from our model, and some suggestion for how modern sensor and computing technology might be used to alleviate problems of self-control

2. Self-control

In its simplest form, the problem of self-control consists of the fact that we tend to choose smaller, sooner rewards rather than larger later rewards *despite* knowing that this is against our best interest. At a descriptive level such situations show a characteristic profile. At T1, a safe distance from the reward, we decide that we prefer the greater reward to be delivered at T3, to the lesser reward delivered at T2. However, at an intermediate time right before T2 we succumb to the imminent lesser reward, which is then followed by regret and lament at T3. Obviously, not all self-control problems are so straightforward, but clearly delineated conflicts between smaller sooner and larger later rewards (what Rachlin, 2000, calls *simple ambivalence*) form the core of the issue of self-control. Importantly, the problem of self-control lies not simply in the act of impulsively choosing an immediate and “lesser” reward, but in doing so *against ones own recognised best interest*. Examples of this would include failure to follow through on decisions to start exercising, or quit smoking, or a constant tendency to put off writing important assignments at school.

But the problem of self-control is not just a problem manifested in the behaviour of certain “weak-willed” individuals: it is a basic, universal and hardwired feature of reward-anticipation in the human brain. A great deal of research has been conducted into choice-behaviour in relation to different temporal distributions of reward, both in animals and humans (see for example Mazur, 2001; Frederick, Loewenstein & O'Donoghue, 2003). The main result that has emerged from this research is a general mathematical function that concisely expresses the diminution of the motivational force of a reward in relation to the length of delay. This is called the *temporal discount function* (Monterosso & Ainslie, 1999). All organisms prefer having immediate access to reward, rather than having to wait for it. The most important empirically derived property of the discount function is its hyperbolic shape (Ainslie, 2001; Monterosso & Ainslie, 1999; see also Read, 2003). What this means is that our subjective evaluation of the reward grows much faster when we are closer to the reward than when we are far from it, and that the mere passage of time can lead to sudden preference reversals.

However, for animals such preference reversals only represent a kind of *manufactured* irrationality. Choice behaviour that produces preference reversals in the laboratory is consistent with models of adaptive foraging in the wild (roughly captured by the adage "a bird in the hand is worth more than ten in the woods"). The motivational systems of rats, pigeons, chimpanzees and other animals are naturally attuned to the reward contingencies of ecologically valid environments, and not to cleverly designed laboratory settings (Rolls, 1999). For humans, on the other hand, temporally induced preference reversals present a serious problem. Unlike most animals, we *constantly* live in a manufactured environment, and the distribution and supply of rewards we face bears little resemblance to the environment in which our species evolved (Duchaine, Cosmides, & Tooby, 2001; but see Kacelnik, 2003). Given our ability to plan ahead and form long-term goals, a natural conflict arises when ancestral systems of reward evaluation entice us with short-term immediate gains. What is revealed by the hyperbolic discounting curve is that failure to follow through on long term goals takes place *just because* we do not have an evolved capacity to wisely, disinterestedly and steadfastly select between short and long-term rewards.

2.1. Computer-Mediated Extrospection

What is it that people do, when they acquire, analyse and act upon self-relevant knowledge? According to folk wisdom, to arrive at such knowledge, people engage in a process of *introspection*, of looking inwards and inspecting the contents of their own minds. Even if this process often is believed to be both fallible and arduous, it is also believed to be more or less transparent to the person involved in the activity. If it is anything in the world that people know with certainty, it is what they themselves think, feel, believe and desire (Goldman, 1993). From this perspective, it would seem that a scheme of capturing and representing aspects of user-context, *for the supposed benefit of the users themselves*, would be of limited value. Such information, it seems, would at best be redundant and superfluous, and at worst a gross mischaracterisation of the user's true state of mind.

On the other hand, common-sense psychology has always acknowledged an imperfect access and command over *some* aspects of our mental lives. Memory, for example, is a case at hand: it is common knowledge that processes of encoding and remembering often are fragile and sometimes inscrutable. The fact that memory is a fundamentally reconstructive process, often at risk of seriously distorting the past, also seems to be generally agreed upon (Buckner & Wheeler, 2001; Paller & Wagner, 2002). Here, then, it is obvious that context aware, memory-enhancing technology could provide a valuable service to users (Mann, 2001; Beigl, 2000). However, computing technology can do more than just emulate the old diary function, and does not have to rest content with capturing information that would have been available to the user if she only had been more attentive or vigilant. Technology can provide information about the state of the user that is *uniquely accessible* by such means. For example, functional magnetic resonance imaging (fMRI) measures of brain activation taken at the time of encoding of a certain material, can (in contrast to the people doing the actual encoding), accurately *predict* levels of recall for a period of several weeks afterwards (Schacter & Dodson, 2001). Similarly, fMRI activity-measures obtained concomitantly with an event can be used to separate out *true* from *false* memories about that event – to a degree not nearly approximated by the remembering agents themselves (Schacter & Dodson, 2001).

Once one starts questioning the scope of our introspective access, it soon becomes clear that it is even more circumscribed than what first appeared to be the case. Memory, for example, is much more than a simple process of routine encoding and retrieval: it is an inseparable component of reasoning and reflection, and deeply involved in our concurrent efforts to gain self-knowledge and regulate our behaviour (LeDoux, 2002). As an illustration, take the case of phobia. For many phobias the subject is unambiguously and acutely aware of the fact that the specific fears they harbour (spiders, open spaces, heights, etc.) are irrational and unreasonable, while at the same time completely failing to act upon this belief (Öhman & Mineka, 2001). The reason for this being that the phobias have been laid down as memories by dedicated subcortical fear-learning mechanisms that are all but cognitively impenetrable (Medina et al., 2002). We simply cannot "look inwards" and divine or correct the workings of these brain systems without extensive training, external prompting, or deliberate relearning (as is practised in cognitive-behavioural therapy and other similar techniques). But phobias are only an extreme case of the constant, day-to-day multilevel learning and responding that takes place in our lives. Using methods of implicit measurement (including everything from hulking basement-dwellers like fMRI scanners, to simple ambulatory sensing of galvanic skin responses) striking dissociations between subjective experience and cognitive/emotional activity have been established in a wide variety of domains and behaviours (e.g. see Gazzaniga, 2000, for a wealth of examples). Evidently, the process of introspection is powerless to survey and regulate a great and important part of our mental economy.

In fact, in our view, these examples demonstrate a general principle about the human cognitive architecture: implicit processing of one or other variety is the norm, not the exception (Rolls, 1999; Dehaene & Naccache, 2000). Even most forms of learning have strong elements of implicitness. The competencies we acquire tend to be anchored in the specific tissues that are modified by training.

They are *embedded* competencies, in the sense that they are incapable of being transported readily to be brought to bear on other problems faced by the individual, or shared with other individuals. *It is knowledge in the system, but not yet knowledge to the system* (Clark & Thornton, 1997).

Importantly, this does not mean that we are powerless to acquire and act upon knowledge about ourselves. Implicit knowledge can be observed in its external, somatic and behavioural manifestations, and it can be subjected to educated, situational "probes" (or sometimes just plain trial and error), in order to generate significant patterns of reaction. We call this process *extrospection*, as we believe it makes for a salient and informative contrast to the traditional concept of *intro*spection. In its basic form, extrospection involves the observation and extraction of regularities that represent the outward expression of implicit information processing in the brain. From such regularities (or sometimes even single instances) the extrospecting agent must infer likely causes and reasons for their occurrence. But, as we mentioned above, extrospection can also involve subtle "provocations" of specialist brain sub-systems, in order to evoke noteworthy reactions (i.e. a sort of quasi-experimental approach to self-reactivity), indicating how external feedback-loops can be used to probe and direct our own brain-internal processes[1].

A currently much discussed example of this latter process is the use of peripheral emotional reactions to guide and constrain decision-making. Studies have shown how loss of peripheral emotional reactions (such as galvanic skin responses) can result in critically impaired decision-making on a variety of tasks (Bechara et al., 2002). The theory behind this being that if we loose extrospective access to the embodied wisdom of our bodily reactions (our so-called somatic markers), then we also loose some of our ability to

1. We specifically want to stress that there can be no clear dividing line between seeing these strategies as unconsciously applying themselves when the situation calls for it, and us deliberately and consciously employing them in the service of a particular goal. As far as cognitive operations go, there is no systematic relation between the complexity of a process and whether it executed in a conscious manner or not (see Dehaene & Naccache, 2000).

make fast and appropriate responses to everyday choice-situations.

Still, as a process of inference, extrospection is subject to the same limitations and problems as any other form of reasoning. In the following sections we elaborate upon possible ways in which the process of extrospection can be augmented by the use of sensor and computing technology – what we call computer-mediated extrospection (CME). CME has many potential uses in the wider process of self-regulation, but here we focus on its particular application to problems of self-control.

2.1.1 Computer-Mediated Extrospection and Self-Control. The starting point for many discussions of self-control is the observation that people are often aware of their self-control problems, but seldom *optimally aware* of the way these problems are expressed in their behaviour, or under what contingencies or in which situations they are most prone to lapses in control (what is called *partial naiveté* in behavioural economics). Most likely, this is due to a mix of biased self-perception, cognitive limitations, and lack of inferential activity (Frederick, Loewenstein & O'Donoghue, 2003). Here CME could serve an important role in correcting faulty self-perceptions. The types of systems we envisage show considerable overlap with initiatives in personal imaging and remembrance agents within wearable and ubiquitous computing (e.g. Mann, 2001; Rhodes, 2000; Singletary & Starner, 2001). However, CME would differentially emphasise the elucidation of information with specific relevance for self-knowledge and self-regulation (not just any task in which augmented memory could be employed). Within this domain, we see four rough categories of CME-tools.

1. *Enhanced Perception*. As a first measure of a CME-tool geared towards improving self-perception, the focus would be on capturing and representing valuable information in our immediate surrounding that we normally fail to register and/or encode, but which we generally believe ourselves to have at least some inkling of. While it may seem like the category of things we falsely believe ourselves to have seen, heard, felt, etc., ought to be very small, evidence suggests otherwise. As the phenomenon of change blindness (i.e. of not noticing potentially gross and remarkable changes to scenes or pictures under conditions of degraded low-level motion information) makes clear, the essence of vision is not as a *form of representation*, but rather a *mode of exploration* (O´Regan & Noe, 2001), and that few things in our near-self environment are registered and retained in any enduring detail (O´Regan & Noe, 2001). On the other hand, our avowed self-knowledge about such matters tends to assume the existence of a much more detailed and reliable impression, and this can lead to a wide variety of self-related misconceptions (Levin et al., 2000).

2. *Macro Prediction.* As a possible means of mitigating problems of self-perception CME could also be used for purposes of *macro prediction*, by finding subtle regularities in behaviour over time and situations. Even if humans are obsessive, incessant and adept pattern-recognisers (whether we know it or not) we are ill suited to process data that is scattered over many different contexts and time-scales. The role of CME could be one of *personal data-mining* (Clarkson, 2002), to discover quirks of acting and responding that are well nigh invisible from the subjective perspective.

3. *Self-Monitoring*. It is also of great importance to apply CME to capture and represent information that we normally successfully access and monitor, but which we sometimes *momentarily* fail to survey. Studies have shown that while humans are quite capable at self-monitoring when given clear directives and timely external prompts, performance quickly deteriorates under natural conditions (Rachlin, 2000). (Compare not trying to scratch an itch under stern scrutiny in the doctor's office, and not scratching it later while watching TV.) The degree of self-monitoring, in turn, greatly influences the nature of our self-control behaviour. There is a big difference between smoking a cigarette that happens to be the 24th of the day, and being aware that one is about to light up the 24th cigarette for the day. The simple fact of providing accurate monitoring of self-control related context has been shown to markedly reduce the incidence of self-control lapses (Rachlin, 2000). The problem is of course that it is almost as difficult to stay

constantly vigilant and attentive to such context as it is to control the behaviour in the first place. This is an area where the use of context aware technology and CME would be of great use.

4. *Micro Prediction*. CME can also be applied in a more direct and intimate manner to measure and influence cognitive and emotional brain activity. A more complex (and potentially more powerful) form of CME would be to apply ubiquitous sensing and computing technology to explicate the relations between different levels of explanation of behaviour. A standard approach in context-aware applications is to build up relevant context from a variety of simple features (or primitives). Combinations of contextual elements like time, location, position, etc. can be used to derive a specific *action*, or *activity* (Gellersen, Schmidt & Beigl, 2002). In explanations of human behaviour the gold-standard of context-abstraction is the intentional level: the level at which we can determine what *purpose* an action has, what it *means*, and what the agents involved *intend*, *desire* and *believe*. There is ample evidence that the human ability to identify intentional states is built up from many semi-autonomous, interdependent processes (detection of self-propelled motion, eye-gaze, joint-attention, etc. see Malle, Moses and Baldwin, 2001). Used in concert, and applied similarly to both oneself and to others, these mechanisms secure the capture of relevant high-level *patterns* in human behaviour, and give our folk-psychology great powers of explanation and prediction (Dennett, 1991). However, such patterns are still abstractions, and leave out much information that could be (and often is) critical to explanations of human behaviour. Here, CME finds several different uses. Most importantly, CME could enable a user to perform various forms of *micro prediction* of her own behaviour. For example, in the fMRI studies of memory encoding described earlier, the subjects involved did not *intend* to forget the material, or held some odd *beliefs* that made forgetting understandable, it was simply the case that the functional-level, brain-based explanation of the process, was *more powerful* than their own self-explanations. This type of prediction does not necessarily have to be based on "in-skull" measurement. As is commonly the case in human-factors studies of error-performance, attention-lapses, and similar

micro behaviours, it could just as easily be based on surface psychophysiology, or even reaction-time performance (Kramer & Weber, 2000; Parson & Hartig, 2000). In a similar way we envision that CME can be used in the context of self-control to set up series of micro predictions of lapse-critical behaviour in the presence of temptation cues, or specific contexts previously associated with relapse.

However, as important as the process of acquiring and processing self-relevant knowledge by CME, is the further use this knowledge is put to in processes of regulation and control. The possibilities of CME as a new *interface* for ourselves go far beyond simple feedback-control. Only detailed experimentation can determine what function – modulating, communicative, explanatory, metacognitive, rewarding, facilitating, distractive, evidential, etc. – that CME might play in any given ubiquitous and context-aware system. In the next section we discuss how the *output* from CME can play a crucial role in instigating and shaping wider processes of motivation and self-control. We also introduce *distributed motivation*, the second of the two general conceptual tools we believe to be important in engaging the problem of self-control.

2.2 Distributed Motivation

As has become evident from our discussion of the nature of the self-control dilemma, and the various means of attaining self-knowledge (whether by our natural senses, or by CME), there is no simple and patented solution to the problem of self-regulation and control. The interesting question is rather what we ordinary folks do when we decide to set out to pursue some lofty goal – to start to exercise on a regular basis, to finally write that film script, to become a less impulsive and irritable person – if we cannot just look inside our minds, exercise our "will," and simply be done with it.

The answer, we believe, is that people cope as best they can with a heterogeneous collection of culturally evolved and personally discovered strategies, skills, tools, tricks and props. We write authoritative lists and schedules, we rely on push and pull from social companions and family members, we rehearse and mull and exhort ourselves with linguistic mantras or potent images of success, and we even set up ceremonial pseudo-contracts (trying in vain to be our own effective enforcing agencies). Often we put salient markers and tracks in the environment to remind us of, and hopefully guide us onto the chosen path, or create elaborate scenes with manifest ambience designed to evoke the right mood or attitude (like listening to soundtracks of old Rocky movies before jogging around the block). We also frequently latch onto role models, seek out formal support groups, try to lock ourselves into wider institutional arrangements (join a very expensive tennis-club with all its affiliated activities), or even hire personal pep coaches. In short, we prod, nudge and twiddle with our fickle minds, and in general try to *distribute* our motivation out into stable social and artefactual structures in the world. Like Odysseus facing the Sirens we often know that we will find ourselves in conditions where we are likely to do something detrimental to our long-term goals, and like Odysseus tying himself to the mast we would often like to be able to *self-bind* or *precommit*, and avoid or resist such temptations.

While various extrospective processes provide the core input and overall shape of our proposed self-regulatory efforts (identifying needs, judging the effectiveness of potential measures, testing solutions to get crucial feedback, etc.) the general strategy of using stable features of the environment to scaffold the process of goal-attainment deserves a special mention. This is what we call *distributed motivation*. As such, distributed motivation is a subclass of the well-established theory of *distributed cognition* (Hutchins, 1995a). Distributed cognition deals with computational processes distributed among agents, artifacts and environments. It is a set of tools and methodologies that allow the researcher to look beyond simple "cognisant" agents, and shift the unit of analysis to wider computational structures (among which the human brain of course is an important part). Distributed motivation aims to achieve the same shift of emphasis in the realm of motivational problems as distributed cognition has done for problems of reasoning (Zhang, 1997), memory (Hutchins, 1995b), and collaboration (Hutchins, 1995a; Rogers & Ellis, 1994). We do not believe there is any principal difference between the "cold" cognitive phenomena normally studied, and the "hot" motivational and emotional processes that are our focus here.

The adoption of an explicit framework of distributed motivation will not only provide a platform in the search for potential remedial applications, but also, we believe, capture often overlooked aspects of how people actually go about trying to overcome problems of self-control (see Elster, 2000, for a similar sentiment). Primary among these aspects, and one of the most central features of our notion of distributed motivation, is the concept of *precommitment* or self-binding.

The tale of Odysseus and the Sirens is a standard illustration of this principle (Elster, 2000; for an in depth treatment, see Sally, 2000a,b). Odysseus, knowing the inevitable consequences of the siren song, orders himself to be tied to the mast (and plugs the ears of the oarsmen crew) thus arranging the environment in such a fashion as to allow him to sail by unharmed. Going back to our outline of the self-control problem, Odysseus suspects at time T1 that he will most likely experience a preference reversal at T2, and so he guarantees, by precommitment, that his original preference will not be violated, and receives the larger long-term reward at T3. What we would like to argue here is that the image of the clever Odysseus foiling the Sirens, might serve as a promising template for the design of modern remedies based on ubiquitous and context-aware technology. While people generally strive to approximate the Odyssean ideal in their daily self-regulation behaviour they seldom manage to create conditions of precommitment stable enough to sustain them through complex and difficult problems. As sure as the fact that the majority of folk-strategies of self-control have been tried and tested in harsh conditions of cultural evolution, or over the full life span of incessantly extrospecting individuals, and that they embody considerable pragmatic wisdom, is also the fact that they fail miserably when looked at on a societal scale (e.g. the extreme pervasiveness of failures to self-regulate that we elaborated upon in the introduction).

2.2.1 Distributed Motivation and Ubiquitous Precommitment Technology. The problem with most folk-strategies is of course that *they do not have enough binding power* (sadly the injunctions are often no stronger than the glue on the back of the post-it notes they are written on). For example, an often-told anecdote in the context of research on self-control is that of the young Afro-American man that made a "powerful" commitment to pay US$ 20 to the Ku Klux Klan every time he smoked a cigarette. In contrast to many other cases it is easy to understand the force this commitment *might* have on his behaviour, but the fact still remains that once he has succumbed to the temptation, nothing really compels him to transfer money to the KKK. But if no such crucial deterrent for future behaviour can be established, then why on earth should he adjust his behaviour in relation to the commitment to begin with? Without going into philosophical niceties, it is easy to see that there is something deeply paradoxical about this kind of self-punishment. Indeed, if one really could exert the type of mental control that effectively *binds* oneself to pay the smoking fee to the KKK, then why not just simply bind oneself not to smoke in the first place?

The main weakness of the strategy employed is the lack of enforcement. The key to improving on the strategy is clearly to increase the binding power of the initial precommitment; in this case, ensuring that the "fine" for smoking is reliably incurred and that lapses are reliably detected. There are in fact several possible solutions, both to monitoring and enforcement. Of the more extreme variety are the agencies that offer round the clock surveillance of dieters and smokers. Although effective, there is a very understandable general resistance to these kinds of schemes. In addition, they are usually cumbersome, inflexible and costly (Ainslie, 1999, 2001; Rachlin, 2000). What is needed are solutions which do not compromise individual integrity, and were the cost of setting up and maintaining the scheme is in parity with the expected benefits. We believe that a ubiquitous infrastructure will be able to meet all of these demands. In the next section, we introduce our schematic model, with illustrations and examples of actual and potential context aware applications to scaffold our self-control behaviour.

3. Self-control scenarios

The issue of self-control is a very complicated phenomenon. Despite the fact that all humans share the same basic cognitive machinery for evaluation of short and long-term rewards (as revealed by the extremely wide applicability of the hyperbolic discount curve) each case has to be evaluated on an individual basis. Obviously, some problems are more severe than others. Common forms of laziness should not be equated with full-blown addiction just because both conditions find their root in similar mechanisms of reward evaluation in the brain. In the following sections we present a model and a discussion of how the conceptual tools we have proposed and discussed in the paper (*computer-mediated extrospection* and *distributed motivation*) can be applied and tailored to the demands of particular self-control problems. We start with comparatively less difficult problems, and move on to harder ones (this progression and our theoretical tools are summarised in figure 1).

3.1 Active Goal Representation

In our discussion of the concept of distributed motivation we catalogued some of the many cultural strategies of self-control that people employ in their daily lives, and noticed how they often fail because of the lack of crucial binding power. However, degree of binding is not the only variable that determines success or failure of any particular attempt at self-control. Sometimes the solution is actually easier than we might first think.

At the most basic level of analysis an often overlooked factor is the nature of the representation of the goals we are striving for. An example from the clinical literature provides a good illustration of this. Patients who have suffered damage to the prefrontal cortex (PFC) often face dramatic impairments in their ability to engage in behaviours that depend on knowledge of a goal and the means to achieve it. They distract too easily, and are said to be "stimulus bound" (Miller, 2000; see also Manuck et al., 2003). Despite this, rehabilitation studies have shown that performance on difficult tasks can be fully restored to the level of control subjects, by the simple use of a wireless, auditory pager system that alerts the patients at random intervals to think about their goals and what they are currently doing (Manly et al., 2002). In this example the pager does not function as a specific memory prosthesis, like a day-planner, or a PDA; it is not telling the patients *what* to do. It is a cheap, global signal that tells them to think about what it was they *really wanted to do*. Similarly, for normal people, there is reason to believe that many of our common failures to follow through on goals and plans, simply stem from an inability to continuously keep our goals active in the face of a bewildering array of distracting (and of course, often tempting) stimuli. Maintenance of behavioural goals is a full time job even for people with perfectly intact prefrontal structures (Miller & Cohen, 2001). "Preferences are not effortlessly stable, the truth is that we manage them, construct them, treat them strategically, we confound them, avoid them, expect change in them and suppress them" (Sally, 2000a, p.690). As is revealed by the wireless pager example, the representational and coordinative power of the human PFC can easily be eclipsed by the intelligence inherent in well-designed cultural artifacts and environments.

Thus, the first tier in any CME-based program for alleviating problems of self-control focuses on maintaining important goals in an active state. Specific types of enhancements to prospective memory exist in countless forms: from post-it notes, to computerised calendars, to ubiquitous context-aware systems like MemoClip (Beigl, 2000) that allow users to associate items or actions to be remembered with specific geographical locations. More general systems, like the wireless pager system described above, have been far less extensively explored. This is unfortunate, because such systems could occupy an important niche that traditional remembrance agents cannot fill. What CME-systems like the wireless pager promise to do, is to act like a *pacemaker for the mind*, a steady signal or beacon to orient our own thinking efforts. It would not require us to specify all our actions in advance (and then give reminders to do those things), but instead encourage us to think back, and apply the knowledge of our prior goals to whatever situation we happen to find ourselves in at the time of the alert.

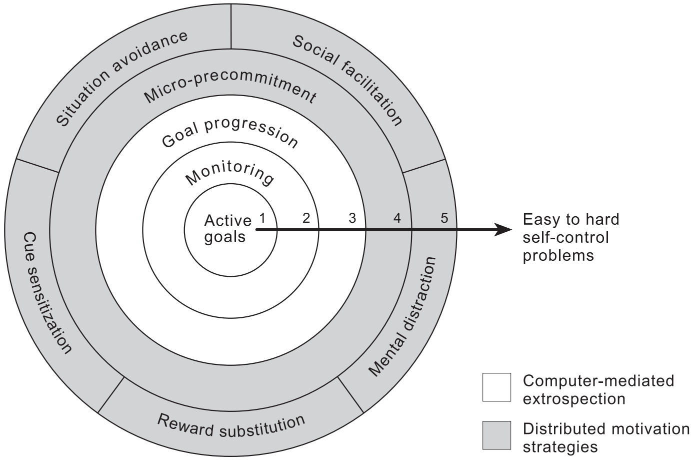

Figure 1. Basic model and applications of the concept of CME and distributed motivation to the problem of self-control. The five circles represent a progression from (comparatively) easy to harder problems. The outer circle contains some representative examples of cultural strategies of distributed motivation that can be plugged into any scheme of precommitment. The fact that the model does not cover other more traditional approaches to self-control (i.e. purely mentalistic approaches, or pharmaceutical interventions) should not be taken as evidence of an opposition to such endeavours; it is only meant to represent approaches that are amenable to manipulation by sensor and computing technology.

A further reason to explore such applications comes from recent findings in basic learning theory. Nelson and Bouton (Nelson & Bouton, 2002; see also Myers & Davis, 2002) have found that a basic asymmetry exists between initial learning in any domain, and subsequent attempts at unlearning such behaviour (for example, eating or drinking habits we would like to change). With few exceptions, initial learning is far less *context-dependent*, while attempts at unlearning generally only work in the specific context where the training took place (for example, in a specific environment, or in a specific state of mind, or even at a specific time, see

Nelson and Bouton, 2002[2]). This means that the risk of relapse is always great unless meticulous care is taken to control for contextual variables that could be of importance. However, Nelson and Bouton (2002) have also shown that this problem can be substantially alleviated by conditioning the retraining to a salient object that is accessible in practically any context (i.e. the object in effect works as a portable context). In the light of the previous discussion, a system like the wireless pager described by Manly et al. (2002) could, with proper preparation, work both as a beacon that is used to re-engage attention to our goals and simultaneously as a signal to (more or less automatically) *inhibit* our bad habits. This would be a powerful example of computing technology that supplies and blends influences from both the "cognitive" and the "motivational" domains.

3.2 Self-Monitoring

The second tier of defence against self-control lapses introduces a more powerful form of CME: one that couples the randomised alert with a context-aware system able to recognise user behaviour that may signal impeding breakdown, or react when it finds the user in specific "contexts of temptation." In our previous discussion of CME we mentioned the fact that, in relation to real-life self-control problems, people often fail to uphold a sufficiently high level of self-monitoring. The phenomenology of lapse behaviour is often completely bereft of any feeling of us having weighed and considered different alternatives, and then finally succumbed to the temptation. Instead we often just *find ourselves*, habitually or absent-mindedly, having performed the act we wanted to avoid.

CME designed to support user self-monitoring could be employed on a scale of both *macro* and *micro*-prediction (see the discussion in section 2.2). Macro prediction would be enabled by sifting through large amounts of context-data relating to lapse-critical behaviour (e.g. neural network approaches, or Bayesian

---

2. Technically, this means that learning to break a bad habit does not involve *unlearning* the old patterns, but rather that *a new form of learning* has been established that (in certain contexts) *inhibits* the old learning. For details, see Nelson and Bouton (2002).

user modelling) and identifying "danger-cues" that could serve to augment and increase the self-knowledge of the user. Micro prediction, on the other hand, would be based on more intimate context measures like the psychophysiological state of the user. In this case, the prediction should be situated at the moment of activity, and come (minutes or seconds) *before* the actual action is performed. For some types of self-control problems this will be comparatively easy. For example, any goals having to do with strong emotions (like trying to become a less aggressive person, or trying to stifle unproductive anger in marital disagreements) will be an ideal target for CME micro prediction. As Elster (2000) has pointed out, advice about emotion regulation most often fail simply because they come *after* the unwanted emotion has already been aroused, and taken full effect upon behaviour. At an earlier stage such advice might have been perfectly effective (i.e. here the proper assessment of the need for self-control is as important as the control itself). Considerable research already exists on psychophysiological markers that indicate the implicit build-up or expression of emotional states, not only for anger and aggression, but also for more subtle conditions like frustration, stress and anxiety (e.g. Caccioppo et al., 2000; Healy & Picard, 1998). Promising efforts are also underway to identify similarly predictive profiles for less obviously emotional behaviour like smoking and gambling (Warren & McDonough, 1999; Blanchard et al., 2000). To increase the chances of finding predictive regularities, CME-technology would add an additional layer to these techniques by allowing the measurements to be individually calibrated over time and multiple contexts. As an example of this, the recently launched *BioMod* project hosted by the MIT Affective Computing Group, aims to develop individually tailored psychophysiological markers of craving-induced stress in smokers trying to quit, and to use this in a large-scale prevention program (more on this in section 3.5 below).

3.3 Goal Progression

Following up on the theme of self-monitoring, the third tier introduces devices or technologies that enable us to better appreciate our level of goal progression. As we mentioned in the earlier discussion of CME, there is a world of difference between lighting up a cigarette that happens to be the 24th of the day, and knowingly and willingly smoking the 24th cigarette of the day. But while CME technology could provide substantial help with monitoring of goals in relation to clear cut objectives like dieting or smoking (it is a relatively straightforward task to implement context-aware devices that could count the amount of calories or cigarettes consumed) it promises to provide an even greater impact in relation to goals that are more abstract, nebulous or distantly long-term. For example, imagine someone that has decided to become a more amiable and caring person. How would she go about fulfilling this goal, and how would she know when she has fulfilled it? One solution that is realisable by means of context-aware technology is to operationalise the goal in such a way as to be able to get discriminating feedback on the outcome of her behaviour. This is a perfect job for context-aware CME-technology. What computers do best is to capture, record, store and analyse data. With the help of ubiquitous or wearable computing devices, conditions of "goal-attainment" could be specified, and used as an objective comparison for the agent involved. Criteria could be set in relation to any behaviour, or activity, or reaction of value that can be automatically captured (number of smiles received, time spend in charity organisation service, galvanic skin responses that indicate deception and lying, environmental contexts that suggest pleasurable social interaction, number of scheduled appointments met in time, etc.). But would this really capture all there is to being an amiable person? No, probably not, but that does not detract from the fact that any change in behaviour in the direction towards such a goal, would be for the better. In our view, the role of CME in such cases could be seen as a form of *scaffolding* that get people started in the direction towards some abstract or long-term goal. When the behavioural change has gained some momentum, the CME-scaffolding can be dropped in order for more complex (and less measurable) behaviours to flourish.

Another similar, but subtly different role for computational technology in monitoring goal-attainment and goal-criteria is provided by Ainslie (2001). He discusses the difficult problem of trying to establish self-controlled behaviour by applying and following *principles*. He argues that in the cultural sphere, and over the lifetime of an individual, a natural evolution of principles takes place, such that (with very few exceptions) principles come to evolve away from what we ideally would like them to do, to instead focus on what is clear and simple and easy to uphold. Thus, an alcoholic that is lucky enough to recover, does not recover as a "social" drinker with a controlled (and presumably) positive intake of alcohol, but as one that abstains from all forms of drinking (Ainslie, 2001; see also discussion in Rachlin, 2000). Total abstinence as a principled approach is much easier to uphold because it leaves no room for subjective interpretation (a beer together with a steak is no real drink, another drink will not hurt me because I have no more cash on me, etc.), and so it does not put the user on a slippery slope. On the other hand, as Ainslie (2001) forcefully argues, what such principles completely ignore, is that this situation might often not be anywhere near what the subject would really want their lives to be like. Again, what CME can bring to this situation is the promise of using computing technology to precisely measure conditions of behaviour and criteria for goal-attainment, in order to effectively emulate the function of principles but without having to settle for the few cases that are so clear cut that our ordinary senses can reliably tell them apart (i.e. we could imagine that with finely tuned sensor and computing equipment, the "social" drinker could live by a CME augmented principle that said that she is only allowed to drink once every other month, or only a certain amount each week, or only if she is at a party of a certain size, etc.).

3.4 Micro Precommitment

Returning now to the core question of time-inconsistent reward evaluation, the fourth tier of defence brings us back to the issue of distributed motivation and methods for self-binding. While active goal representation, swift and accurate self-monitoring, and monitoring of goal-progression are important CME-strategies, they are clearly less applicable in cases of genuine reward conflict. In such cases, precommitment is the right strategy to apply. On the other hand, reward-conflicts come in many different flavours, and often it is not the binding power as such that determines the value of any specific scheme of precommitment. As we outlined in our earlier discussion of ubiquitous precommitment technology, what technology has to offer the age-old strategy of precommitment (apart from more binding-power) is a much-lowered cost and a much-increased range of operation. This is good news, because some species of precommitment need to be fast and easy to set up, and should come at a very low cost. For example, we have remote controls for many electrical appliances that enable us to turn them on and off at our convenience. But we have no remotes that allow us to turn appliances off in a way that, within a set limit of time, *we cannot turn them on again* (for TV and web-surfing, we have things like parental or employer control devices, that can block certain channels or domains, but we have no effective equipment for *self*-binding). We can of course always climb under the sofa, pull the plug and the antenna from the TV, and put them in a place we cannot easily reach (to make TV-viewing relatively inaccessible), but such ad-hoc manoeuvres are generally too costly and cumbersome to perform in the long run. The trick is to strike a balance between inaccessibility and flexibility. That is, for many behaviours and situations we would like to be able to make quick, easy, but transient precommitments, that allow us to move beyond some momentary temptation, but then expire so as not to further limit our range of alternatives. We call this *micro precommitment* (MPC). MPC finds its primary use when the temptations we are dealing with are not overwhelming, but still noticeable enough to bring us to the fall.

As an example, imagine a cell-phone based location-aware system (using GPS or any other modern positioning technique) where we can instantaneously "tag" different places from which we wish to be kept. The mechanism for tagging could be as simple as having the phone in the same "cell" as the object to be tagged, or having a place-map database in the phone that allows for distance independent blocking. Let us now say we have a minor shoe-shopping compulsion, and walk around town on an important errand. Walking down the street with this system we could, with just a brief moment of forethought, tag an upcoming tempting shoe-store. The tagging could have any number of consequences, like locking our wallet or credit-card, or even tuning the store-alarm to go off if we enter the premises. The point of MPC is *not* to set up consequences that represent maximally strong deterrents. Quite the opposite: it is a technique suited for temporarily bringing us past small but nagging distractions. Tomorrow, when we have no important errands anymore, we might want to shop for shoes again, and would not want to spend our time unwinding a too forceful and elaborate precommitment scheme. In fact, since MPCs, in our view, should be as easy and cheap as possible to instigate, they *should also not be allowed to have costly or long-term consequences*.

3.5 Precommitment

The final tier in our program starts out where MPC leaves off. While MPCs are swift and cheap and play with low stakes and short-term consequences, regular precommitment holds no such limits. For precommitment the amount of binding power and the cost of engagement are determined in relation to the magnitude of the problem, and may be as strong as any agent desires. In contrast to MPC, regular precommitment should *not* come easy. To make sure that the binding represents a "true" preference, a certain amount of inertia ought to be built into any precommitment decision procedure (for a sensitive discussion of how to handle this problem, see Elster, 2000). For example, some larger casinos give patrons prone to too much gambling the option of having themselves banned from playing. Since casinos are generally equipped with rigorous security and surveillance systems, the ban can be very effectively enforced. However, one can not just walk up to the entrance cashier and ask to be banned. The decision must be made in dialogue and with council from the casino management, because once you are banned *the casino will not be coaxed into letting you in again*. As would be expected from a compulsive gambler, you soon find yourselves back at the gates trying to undo your former decision. It is at this point that the casino enforces the bind by bluntly disregarding your pleas (and if the commitment was made in too light a manner, this would be an unfortunate outcome).

As we explained in our earlier discussion of ubiquitous precommitment technology, the prime strength of such technology is the manifold of new possibilities for manipulating varieties and degrees of binding it introduces. The question is: are these benefits substantial enough to allow us to fashion realistic scenarios for the alleviation of more difficult problems of self-control, such as craving and addiction? We believe so.

Craving and addiction are extremely difficult topics to approach. Behavioural abnormalities associated with addiction are exceptionally long-lived, and currently no reliable remedies exist for the pathological changes in brain-reward systems that are associated with prolonged substance abuse (Nestler, 2001; Everitt, Dickinson & Robbins, 2001; Robinson & Berridge, 2003). With reference to precommitment, it is sometimes said that it is a limited strategy for handling things like addiction, because in the addicted state we supposedly never find a clear *preference platform* from which to initiate the precommitment (i.e. we do not know which of our preferences that are the "true" ones). Rachlin (2000) writes: "Instead of clearly defined points of time where one strong preference gives way to its opposite we generally experience a continuous opposition of forces and apparently random alternation between making and breaking our resolutions" (p. 54). This state of *complex ambivalence* (as Rachlin calls it) also makes it likely that a fierce arms-race will be put in motion by the introduction of any scheme of precommitment, where the addicted subject will waste precious resources and energy trying to slip through the bind of the commitment. The drug Antabuse illustrates these problems. If you take Antabuse and then have a drink, you will experience severe pain. Thus, taking Antabuse is a form of precommitment not to drink alcohol. However, alcoholics have been known to subvert the effects of the drug by sipping the alcohol excruciatingly slowly, and some even drink the alcohol despite the severe pain (Rachlin, 2000). Also, the outcome of Antabuse treatment has been generally less than satisfying because many alcoholics decide against taking the drug in the first place.

In our view, this example should be taken as a cautionary tale for any overly optimistic outlook on the prospects of precommitment technology to handle really tough cases like addiction, but we do not believe it warrants a general doubt about our approach. As is evident by the fantastically prosperous industry for the supply of services and products that purports to alleviate problems of self-control (in practically any domain of life) people are willing to take on substantial commitments, in terms of time, energy, and resources, to change their current ways of life.

Take smoking as an example. What would a ubiquitous precommitment scheme for helping smokers to quit look like? Firstly, as a foundation, some means of detecting the presence or absence of smoking-related context is needed. The context could be built from observation of the actual smoking, from traces of smoking (from smoking-related behaviour patterns, or from psychophysiological concomitants of smoking), and many types of sensors could be used to generate the match. For example, one sensor-platform that might be used in the near future to provide robust and efficient measurement, is in-blood substance detection. In relation to diabetes treatment, Tamada, Lesho and Tierney (2002) describe a host of emerging *transdermal* (through the skin) techniques for measuring glucose levels in the blood[3]. While not perfected yet, such sensors can be worn continually and unobtrusively by diabetics to efficiently monitor and manage their blood sugar levels. A similar system could easily be envisaged for nicotine[4]. Yet, as Gellersen, Schmidt, and Beigl (2002) have shown, a combination of many cheap and overlapping environmental sensors (i.e. things like temperature, acceleration, light, movement, etc.) might provide equally robust context-measurement as a specialised subcutaneous device.

The great boon of ubiquitous precommitment technology is that once the basic sensing of context is in place (in the previous fictional example, transdermal nicotine blood level detection), a multitude of distributed motivational strategies can be latched onto it, and varieties of binding can be added or subtracted

---

3. Nicotine delivery skin patches are an example of transdermal technology working in the other direction, where the molecule of interest is moving into the body rather than out of it.
4. If we want to limit ourselves to existing technologies, CO is considered to be a very reliable indicator of smoking, and products monitoring the CO level in exhaled air have been used for a number of years (e.g. the Smokerlyzer™). Using saliva samples is currently the fastest and least obtrusive way of detecting nicotine, and products for this purpose have also been around for some time (e.g. NicAlert™, Accutest®).

depending on the nature and severity of the case. The versatility of the platform also allows for overlapping and partially redundant incentives to be put in place. To take a dramatic example, for providing strong and relentless binding, a wireless bracelet for nicotine monitoring could be hooked up directly to the bank account of the participating subject, and simply withdraw money in proportion to the amount of smoking the subject does. But to prevent loss of money, an anticipatory CME backup-system that detects "lapse-critical" behaviour (as described in section 3.2 above) could be employed alongside the nicotine-bracelet, and make automatic support calls to other participants in the program if the subject is in danger of taking a smoke; a very similar approach to this is taken in the MIT BioMod project we described earlier. The extracted psychophysiological markers of "lapse-critical" stress levels will be used to automatically relay cell-phone calls to a support centre where trained professional can answer to the needs of the subject. In all, we foresee that while exceptionally strong single precommitment criteria can be put in place (i.e. you loose all your money if you smoke one single cigarette), it is the possibility of mixing and merging many less forceful strategies in one system that will provide the greatest benefits. Most likely, venerable cultural strategies like situation avoidance (e.g. the shoe-store "tagging" example), social facilitation, reward-substitution, etc., will experience a strong resurgence in the hand of ubiquitous technology for distributed motivation.

4. Discussion

4.1 Flexibility and Rigidity

Some researchers have expressed great pessimism about the ability of context-aware systems to make meaningful inferences about important human social and emotional states, and believe that context aware applications can only supplant human initiative in the most carefully proscribed situations (Bellotti & Edwards, 2001). We are in no position to assess the finer details behind this pessimism, but it must be noted that this problem is far less pressing for the proposed domain of ubiquitous self-control technology discussed in this article. Precommitment technologies offer people the option of temporary, but forceful, binding, aided by computer systems that will not be swayed or cajoled, and *it is through their very inflexibility that these systems have the potential to support individual self realisation*. As Dennett (2003) notes, in the domain of self-control effectively constraining our options gives us more *freedom* than we otherwise would have had.

Nevertheless, the rigidity of these technologies may sometimes be a weakness. Although precommitment technology increases the likelihood of attaining individually set goals, there is the attendant risk that people will lock themselves into inappropriate precommitments, and waste time and effort fulfilling needless obligations. It can be relatively easy to precommit, but comparatively hard to foresee potential conflicts with other valued goals and preferences one might have. Preferences are also subject to change and the future can bring unexpected opportunities, as well as emergencies, that we want to be able to respond to. In addition, there is the possibility that, once precommitted, some people will expend wasteful resources on increasingly elaborate countermeasures. These concerns must all be taken seriously, but are not as severe as they might appear at first blush. Clearly, the proposed systems will have to leave room for a host of pre-programmed contingencies, as well as a fixed number of predetermined digressions. If precommitment does come with a price (temporally limited freedom) this must be taken in relation to the valuation of the goal one wishes to attain; if a particular precommitment seems arduous this also has to be judged in relation to potential benefits. Ultimately, people will be free to use these systems or not, as they see fit, and to weigh potential benefits against possible costs.

A degree of inflexibility is essential to the successful working of these kinds of systems, but as we discussed earlier, we foresee a range of different kinds of binding, of various degree and type, that can be combined in regimes suitably coupled to particular issues. In a possible scheme, the range of permissible actions is large to start with, but slowly curtailed in response to flagging willpower. In an alternative scheme, permissible actions are limited at the onset, but then expand as the need for support slowly wanes (this kind of regime might be seen as a form of motivational scaffolding or training wheel).

4.2 Ethical Considerations

With the introduction of new technology also come new ethical considerations. With lots of information about the user being picked up and circulated (information about location, behaviour, affective and cognitive states etc.) there is the risk that the information could be put to unsavoury use. A widely shared worry is that this kind of information is a threat to privacy. The problem of privacy is one that besets the whole field of ubiquitous computing, and there have been some viable and thoughtful suggestions of how this could be handled (Dey, Abowd & Salber, 2001; Bellotti & Edwards, 2001). Nevertheless, with more intimate, psychological, measures afloat the problem of privacy is perhaps even more pressing.

A standard solution to the problem of privacy in context-awareness, is to increase the "transparency" of the applications, making users aware of what kind of information about them is being registered, and what actions are about to be taken in response. On the other hand the obvious problem with trying to increase the transparency of contextually aware applications is that constant requirement of notice could easily overwhelm users, and disrupt their activities (think of such a relatively simple task as management of browser cookies). It is a reasonable question to ask whether users would provide attention and direction if they were constantly bombarded by requests from all kinds of systems (temperature and light settings, image-capture, notes and file-sharing, driver safety customisations, etc.). Studies have shown that people are both poor at handling such updates, and unwilling to receive them (Belloti & Edwards, 2001; Ackerman, Darrell & Weitzner, 2001). In this regard, ubiquitous CME and precommitment technology have a clear advantage over many other context-aware applications. Again, what we would like to stress here is that these concerns are not nearly as pressing for a scheme of representing and augmenting user-perceptions of context – *to the users themselves*. The emphasis on explicitness of interaction (at the loss of some ease and efficiency) is not a problem for the manufacturing of self-control technology; quite the opposite. Given the personal importance potentially at stake in such examples it would be dangerous and irresponsible to allow the process to proceed entirely implicitly. People already spend a great deal of time and effort trying to regulate and manage their cognitions, emotions and behaviour. Our project only proposes to usurp resources already devoted with scant success to similar causes.

Another ethical concern for the prospect of ubiquitous CME and precommitment technology, is that with the availability of these kinds of systems there is a risk that people will be put under undue pressure to employ them (by family members or employers, or maybe even government agencies). If precommitment is too easy to set up, and the binding forceful, there is the risk that people get stuck in precommitments they wouldn't have chosen "under a calm moment of reflection." We must therefore ensure that precommitments are not entered into under duress, but at appropriate times and for appropriate reasons. A related concern is that pressure might come from the system itself. Systems like these could be purposefully designed to be persuasive: to lure users into setting up various kinds of precommitments and obligations. This is one point at which we clearly differ from the avowed, but related, goals of *persuasive computing* (e.g. Fogg, 2003). Although we have selected some possible societally beneficial areas for remedial action, our intent is not one of persuading people to participate. All along, the premise of our work have been that self-control problems only apply to situations in which the subject herself considers it to be a problem (choosing a "lesser" reward *against ones own recognised best interest*). This does not mean that a "persuasive" or paternalistic stance is never justified (see discussion in Fogg, 2003; and O´Donoghue & Rabin, 2003), but it has not been part of our concern here.

4.3. Summary

We have provided a basic model and a host of examples of how the twin concepts of CME and distributed motivation can be applied and tailored to problems of self-control. Our scheme of classification is intended to provide an overview of the impact ubiquitous sensor and computing technology might have on the domain of self-control. It is our hope that the model and our discussion will provide a principled and useful way for designers of human–computer interfaces and context-aware systems to approach the domain of self-control, as well as to provoke further debate on the possible role of computing technology in matters of human motivation.

The technologies and theories proposed here are, we believe, well grounded, but need to be tested in an arena of real self-control problems and against a background of technological constraints. For the future, we envision precommitment technologies and tools of computer-mediated extrospection that can be configured by the users themselves, in ways and for purposes we cannot yet anticipate. It is in the ecology of devices, human needs and ingenuity that the field will take shape.

Acknowledgements

We would like to thank the participants of the Lund University Cognitive Science seminar for their input on an earlier draft of the paper. Special thanks go to Jens Månsson for his contributions at the early stages of this project.

REFERENCES

Ackerman, M., Darrell, T., & Weitzner, D. J. (2001). Privacy in context. *Human–Computer Interaction, 16,* pp. 167–176.

Agre, P., & Shrager, J. (1990). Routine evolution as the microgenetic basis of skill acquisition. *Proceedings of the 12 th Annual Conference of the Cognitive Science Society* (pp. 694–701). Cambridge, MA: Lawrence Erlbaum.

Ainslie, G. (1999). The dangers of willpower. In J. Elster and O-J. Skog (Eds.), *Getting hooked: Rationality and addiction* (pp. 65–92). Cambridge: Cambridge University Press.

Ainslie, G. (2001). *Breakdown of will.* Cambridge: Cambridge University Press.

Akrich, M. (1991). The de-scription of technical objects. In W. E. Bijker and J. Law (Eds.), *Shaping technology/Building society: Studies in sociotechnical change* (pp. 205–224). Cambridge, MA: MIT Press.

Appadurai, A. (1986). Introduction: commodities and the politics of value. In A. Appadurai (Ed.), *The social life of things* (pp. 3–63). Cambridge: Cambridge University Press.

Baddeley, A. (1997). *Human memory: Theory and practice, revised edition.* Hove: Psychology Press.

Ballard, D. (1991). Animate vision. *Artificial Intelligence, 48(1),* pp. 57–86.

Barker, R. (1968). *Ecological psychology: Concepts and methods for studying the environment of human behavior.* Stanford: Stanford University Press.

Basalla, G. (1988). *The evolution of technology.* Cambridge: Cambridge University Press.

Beach, K. (1988). The role of external mnemonic symbols in acquiring an occupation. In M. M. Gruneberg, P. E. Morris and R. N. Sykes (Eds.), *Practical aspects of memory: Current research and issues, vol. 1* (pp. 342–346). Chichester: Wiley.

Bechara, A., Ranel, D., & Damasio, A. R. (2002). The somatic marker hypothesis and decision-making. In F. Boller and J. Grafman (Eds.), *The handbook of neuropsychology, vol. 7* (pp. 117–143). Amsterdam: Elsevier Science.

Beigl, M. (2000). MemoClip: A location-based remembrance appliance. *Personal Technologies, 4,* pp. 230–233.

Bellotti, V. M., & Edwards, W. K. (2001). Intelligibility and accountability: human considerations in context aware system. *Human–Computer Interaction, 16,* pp. 193–212.

Bijker, W. E. (1994). Sociohistorical technology studies. In S. Jasanoff, G. E. Markle, J. C. Petersen and T. Pinch (Eds.), *Handbook of science and technology studies* (pp. 229–256). Thousand Oaks: Sage Publications.

Blackaby, J. R. (1986). How the workbench changed the nature of work. *American heritage of invention and technology, fall issue,* pp. 26–30.

Blanchard, E. B., Wulfert, E., Freidenberg, B. M., & Malta, L. S. (2000). Psychophysiological assessment of compulsive gamblers' arousal to gambling cues: A pilot study. *Applied Psychophysiology and Biofeedback, 25(3),* pp. 155–165.

Bleed, P. (1986). The optimal design of hunting weapons: Maintainability or reliability. *American Antiquity, 51(4),* pp. 737–747.

Brinswell, A. (2001). The mind of a terra cotta warrior: the transubstantiation of matter into thought. *Clinamen, 23(1),* pp. 19–69.

Brooks, R. A. (1991a). Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence* (pp. 569–595). Morgan Kauffman.

Brooks, R. A. (1991b). Intelligence without representation. *Artificial Intelligence, 47(1–3),* pp. 139–159.

Buckner, R. L., & Wheeler, M. E. (2001). The cognitive neuroscience of remembering. *Nature Reviews Neuroscience, 2,* pp. 624–634.

Bærentsen, K. B. (1989). Mennesker og maskiner. In M. Hedegaard, V. R. Hansen and S. Thyssen (Eds.), *Et virksomt liv: Udforskning af virksomhedsteoriens praksis* (pp. 142–187). Aarhus: Aarhus Universitetsforlag.

Cacioppo, J. T., Berntson, G. G., Sheridan, J. F., & McClintock, M. K. (2000). Multilevel integrative analyses of human behavior: social neuroscience and the complementing nature of social and biological approaches. *Psychological Bulletin, 126(6),* pp. 829–843.

Caroll, J. M. (1991). Introduction: the Kittle house manifesto. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human–computer interface* (pp. 1–16), Cambridge: Cambridge University Press.

Carroll, J. M., Kellogg, W. A., & Rosson, M. B. (1991). The task–artifact cycle. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human–computer interface* (pp. 74–102). Cambridge: Cambridge University Press.

Clark, A. (1997). *Being there: putting brain, body, and world together again.* Cambridge, MA: MIT Press.

Clark, A. (1998). Where brain, body and world collide. *Daedalus, 127(2),* pp. 257–280.

Clark, A. (1999a). An embodied cognitive science? *Trends in Cognitive Sciences, 3(9),* pp. 345–351.

Clark, A. (1999b). Embodied, situated, and distributed cognition. In W. Bechtel and G. Graham (Eds.), *A companion to cognitive science* (pp. 506–517). Oxford: Blackwell.

Clark, A. (2001). Reasons, robots and the extended mind. *Mind and Language, 16(2),* pp. 121–145.

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

Clark, A., & Thornton, C. (1997). Trading spaces: Computation, representation, and the limits of uniformed learning. *Behavioral and Brain Sciences, 20,* pp. 57–90.

Clarkson, B. (2002). *Life patterns: structure from wearable sensors.* Unpublished doctoral dissertation, MIT Media Lab.

Cole, M. (1991). Conclusion. In L. B. Resnick, J. M. Levine and S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 398–417). Washington DC: American Psychological association.

Cole, M. & Engestöm, Y. (1993). A cultural-historical approach to distributed cognition. In G. Salomon (Ed.), *Distributed cognitions – Psychological and educational considerations* (pp. 1–46). Cambridge: Cambridge University Press.

Cole, M., & Griffin, P. (1980). Cultural amplifiers reconsidered. In D. R. Olson (Ed.), *The social foundations of language and thought* (pp. 343–364). New York: W. W. Norton & Company.

Cole, M. & Wertsch, J. V. (n.d./1999). Beyond the individual-social antimony in discussions of Piaget and Vygotsky. [URL] http://www.massey.ac.nz/~ALock/virtual/colevyg.htm

Dehaene, S., & Naccache, L. (2000). Toward a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79,* pp. 1–37.

Dennett, D. C. (1991). Real patterns. *Journal of Philosophy, 89,* pp. 27–51.

Dennett, D. C. (2003). *Freedom evolves.* London: Allen Lane.

Dey, A. K., Abowd, G. D., & Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human–Computer Interaction, 16,* pp. 167–176.

Diaper, D. (Ed.). (1989). *Task analysis for human–computer* inter*action.* Chichester: Ellis Horwood.

Draper, S. W. (1993). The notion of task in HCI. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel and W. White (Eds.), *Bridges between worlds, INTERCHI'93 Conference proceedings, adjunct proceedings* (pp. 207–208). Reading, MA: Addison-Wesley.

Duchaine, B., Cosmides, L. & Tooby, J. (2001). Evolutionary psychology and the brain. *Current Opinion in Neurobiology, 11,* pp. 225–230.

Dupuy, T. N. (1984). *The evolution of weapons and warfare.* Cambridge; New York: DaCapo Press.

Elster, J. (2000). *Ulysses unbound.* Cambridge: Cambridge University Press.

Engeström, Y. (1999). Activity theory and individual and social transformation. In Y. Engeström, R. Miettinen and R.-L. Punamaki (Eds.), *Perspectives in activity theory* (pp. 19–38). New York: Cambridge University Press.

Everitt, B. J., Dickinson, A., & Robbins, T. W. (2001). The neuropsychological basis of addictive behavior. *Brain Research Reviews, 36,* pp. 129–138.

Fogg, B. J. (1999). Persuasive technologies. *Communications of the ACM, 42(5),* pp. 27–29.

Fogg, B. J. (2000). Persuasive technologies and netsmart devices. In E. Bergman (Ed.), *Information appliances and beyond: Interaction design for consumer products* (pp. 336–360). San Francisco: Morgan Kaufmann.

Fogg, B. J. (2003). *Persuasive technology: Using computers to change what we think and do.* San Francisco: Morgan Kaufmann Publishers.

Frederick, S., Loewenstein, G. & O'Donoghue, T. (2003). Time discounting and time preference: A critical review. In G. Loewenstein, D. Read and R. Baumester (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 13–86). New York: Russell Sage Foundation.

Garbis, C. (2002). The cognitive use of artifacts in cooperative process management: Rescue coordination and underground line control. *Linköping Studies in Arts and Science, 258.*

Gardner, G., & Halweil, B. (2000). *Underfed and overfed: The global epidemic of malnutrition.* Washington DC: Worldwatch Paper, 150.

Gatewood, J. (1985). Actions speak louder than words. In J. Dougherty (Ed.), *Directions in cognitive anthropology* (pp. 199–220). Urbana: University of Illinois Press.

Gaver, W. W. (1991). Technology affordances. In S. Robertson, G. Olson and J. Olson (Eds.), *Proceedings of CHI* (pp. 79–84). New York: ACM.

Gazzaniga, M. S. (Ed.), (2000). *The new cognitive neurosciences, 2nd edition.* Cambridge, MA: MIT Press.

Gedenryd H. (1998). How designers work: Making sense of authentic cognitive activities. *Lund University Cognitive Studies, 75.*

Gellersen, H. W., Schmidt, A. & Beigl, M. (2002). Multi-sensor context-awareness in mobile devices and smart artifacts. *Mobile Networks and Applications, 7,* pp. 341–351.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Goldman, A. I. (1993). Consciousness, folk-psychology, and cognitive science. *Consciousness and Cognition, 6,* pp. 364–382.

Gosden, C. & Marshall, Y. (1999). The cultural biography of objects. *World Archaeology, 31 (2),* pp. 169–178.

Gould, S. J. & Vrba, E. (1982). Exaptation – a missing term in the science of form. *Paleobiology, 8(1),* pp. 4–15.

Govoroff, N. C. (1993). The hunter and his gun in Haute-Provence. In P. Lemonnier (Ed.), *Technological choices* (pp. 227–237). London and New York: Routledge.

Green, T. R. G. & Petre, M. (1996). Usability analysis of visual programming environments: A "cognitive dimensions" framework. *Journal of Visual Languages and Computing, 7(2),* pp. 131–174

Hackos, J. T. & Redish, J. C. (1989). *User and task analysis for interface design.* New York: Wiley.

Halverson, C. A. (1995). Inside the cognitive workplace: New technology and air traffic control. Ph.D. Thesis, Dept. of Cognitive Science, University of California, San Diego.

Hammond, K. J. (1990). Learning and enforcement: Stabilizing environments to facilitate activity. In B. W. Porter and R. J. Mooney (Eds.), *Proceedings of the seventh international conference on machine learning* (pp. 204–210). San Mateo, CA: Morgan Kaufman.

Hammond, K. J., Converse, T. M., & Grass, J. W. (1995). The stabilization of environments. *Artificial Intelligence, 72,* pp. 305–327.

Hazlehurst, B. L. (1994). *Fishing for cognition: An ethnography of fishing practice in a community on the west coast of Sweden.* University of California, San Diego: Unpublished dissertation.

Healey, J., & Picard, R. (1998). *Digital processing of affective signals*. MIT: Technical Report.

Hutchins, E. (1986). Mediation and automatization, *Quarterly Newsletter of the Laboratory of Comparative Human Cognition, 8(2),* pp. 47–58. Reprinted in M. Cole, Y. Engeström and O. Vasquez (Eds.), *Mind, culture, and activity: Seminal papers from the Laboratory of Comparative Human Cognition* (pp. 338–353). Cambridge: Cambridge University Press, 1997.

Hutchins, E. (1990). The technology of team navigation. In J. Galegher, R. Kraut and C. Egido (Eds.), *Intellectual teamwork: Social and technical bases of collaborative work* (pp. 191–220). Hillsdale: Lawrence Erlbaum Associates.

Hutchins, E. (1995a). *Cognition in the wild.* Cambridge, MA: MIT Press.

Hutchins, E. (1995b). How a cockpit remembers its speed. *Cognitive Science, 19,* pp. 265–288.

Hutchins, E. (1999). Cognitive artifacts. In R. A. Wilson and F. C. Keil (Eds.), *The MIT encyclopaedia of the cognitive sciences* (pp. 126–128). Cambridge MA: MIT Press.

Hutchins, E. (2001). Cognition, distributed. In *International encyclopaedia of the social and behavioral sciences* (pp. 2068–2072). Pergamon Press.

Hutchins, E. (2002). Material anchors for conceptual blends. In A. Hougaard and S. N. Lund (Eds.), *Proceedings of The Way We Think symposium on conceptual blending* (pp. 85–111). Odense, Denmark: University of Southern Denmark.

Hutchins, E., & Klausen, T. (1996). Distributed cognition in an airline cockpit. In Y. Engeström and D. Middeleton (Eds.), *Cognition and communication at work* (pp. 15–34.). Cambridge: Cambridge University Press.

Kacelnik, A. (2003). The evolution of patience. In G. Loewenstein, D. Read and R. Baumester (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 115–138). New York: Russell Sage Foundation.

Kaptelinin, V. (1996). Distribution of cognition between minds and artifacts: Augmentation or mediation? *AI and Society, 10,* pp. 15–25.

Keller, C. M. (2001). Thought and production: Insights of the practitioner. In M. B. Schiffer (Ed.), Anthropological perspectives on technology (pp. 33–45). Albuquerque: University of New Mexico Press.

Keller, C. M., & J. D., Keller (1993). Thinking and acting with iron. In S. Chaiklin and J. Lave (Eds.), *Understanding practice: Perspectives on acting and context* (pp. 125–143). Cambridge: Cambridge University Press.

Keller, C. M., & Keller, J. D. (1996). *Cognition and tool use: The blacksmith at work*. Cambridge: Cambridge University Press.

Kirlik, A. (1998). Everyday life environments. In W. Bechtel and G. Graham (Eds.), *A companion to cognitive science* (pp. 702–712). Oxford: Blackwell.

Kirsh, D. (1995a). Complementary strategies: Why we use our hands when we think. In J. D. Moore and J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 212–217). Hillsdale, NJ: Lawrence Erlbaum.

Kirsh, D. (1995b). The intelligent use of space. *Artificial Intelligence, 73,* pp. 31–68.

Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior, 4, 3/4*, pp. 415–452.

Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science, 18,* pp. 513–549.

Kopytoff, I. (1986). The cultural biography of things: commoditization as process. In A. Appadurai (Ed.), *The social life of things* (pp. 64–91). Cambridge: Cambridge University Press.

Kramer, A. F. & Weber, T. (2000). Applications of psychophysiology to human factors. In J. T. Cacioppo, L. G. Tassinary and G. G. Berntson (Eds.), *Handbook of psychophysiology* (pp. 794–814). Cambridge: Cambridge University Press.

Landauer, T. K. (1991). Let's get real: A position paper on the role of cognitive psychology in the design of humanly useful and usable systems. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human–computer interface* (pp. 60–74), Cambridge: Cambridge University Press.

Latour, B. (1986). Visualization and cognition: Thinking with eyes and hands. *Knowledge and Society: Studies in the Sociology of Culture Past and Present,* 6, pp. 1–40.

Latour, B. (1988). Mixing humans and non-humans together: the sociology of a door-closer. *Social Problems, 35(3),* pp. 298–310.

Latour, B. (1991a). Technology is society made durable. In J. Law, (Ed.), *A sociology of monsters: Essays on power, technology and domination* (pp. 103–131). London: Routledge.

Latour, B. (1991b). Where are the missing masses? The sociology of a few mundane artifacts. In W. E. Bijker and J. Law (Eds.), *Shaping technology/Building society: Studies in sociotechnical change* (pp. 225–258). Cambridge, MA: MIT Press.

Latour, B. (1999). *Pandora's hope: essays on the reality of science studies.* Cambridge, MA: Harvard University Press.

Lave, J. (1988). *Cognition in practice: mind, mathematics and culture in everyday life.* Cambridge: Cambridge University Press.

Lave, J., Murtaugh, M., & de la Rocha, O. (1984). The dialectic of arithmetic in grocery shopping. In B. Rogoff and J. Lave (Eds.), *Everyday cognition: its development in social context* (pp. 67–94). Harvard: Harvard University Press.

LeDoux, J. (2002). *The synaptic self.* New York: Viking.

Lemonnier, P. (1986). The study of material culture today: toward an anthropology of technical systems. *Journal of Anthropological Archaeology, 5,* pp. 147–186.

Lemonnier, P. (1992). Elements for an anthropology of technology. *Anthropological Papers, Museum of Anthropology, University of Michigan, 88.*

de Léon, D. (1999). Building thought into things. *Proceedings of the 3rd European Conference on Cognitive Science.* pp 37–47.

de Léon, D. (2002). Cognitive task transformations. *Cognitive Systems Research, 3(3),* pp. 449–459.

de Léon, D. (2003a). The cognitive biographies of things. *Lund University Cognitive Studies, 103.*

de Léon, D. (2003b). Action, artefacts and cognition: A study of people cooking. *Lund University Cognitive Studies, 104.*

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition, 7,* pp. 397–412.

Mahias, M. -C. (1993). Pottery techniques in India: Technical variants and social choice. In P. Lemonnier (Ed.), *Technological choices* (pp. 157–180). London and New York: Routledge.

Malle, B. F., Moses, L. J., & Baldwin, D. A. (Eds.), (2001). *Intentions and intentionality: Foundations of social cognition.* Cambridge, MA: MIT Press.

Manly, T., Hawkins, K., Evans, J., Woldt, K., & Robertson, I. (2002). Rehabilitation of executive function: Facilitation of effective goal management on complex tasks using periodic auditory alerts. *Neuropsychologia, 40(3),* pp. 271–281.

Mann, S. (2001). *Intelligent image processing.* New York, Chichester: Wiley.

Manuck, S. B., Flory, J. D., Muldoon, M. F., & Ferrell, R. E. (2003). A neurobiology of intertemporal choice. In G. Loewenstein, D. Read and R. Baumester (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 139–172). New York: Russell Sage Foundation.

Marr, D. (1982). *Vision.* San Francisco, CA: Freeman.

Mazur, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review, 108(1),* pp. 96–112.

McClamrock, R. (1995). *Existential cognition.* Chicago: The University of Chicago Press.

Medina, J. F., Repa, J. C., Mauk, M. D., & LeDeoux, J. E. (2002). Parallels between cerebellum- and amygdala-dependent conditioning. *Nature Reviews Neuroscience, 3(2),* pp. 122–131.

Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience, 1,* pp 59–65.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* pp. 167–202.

Mokdad, A. H., Bowman, B. A., & Ford, E. S. (2001). The continuing epidemic of obesity and diabetes in the United States. *JAMA 286, 10,* pp. 1195–1200.

Monterosso, J., & Ainslie, G. (1999). Beyond discounting: Possible experimental models of impulse control. *Psychopharmacology, 146,* pp. 339–347.

Myers, K. M., & Davis, M. (2002). Behavioral and neural analysis of extinction. *Neuron, 36(4),* pp. 567–84.

Nardi, B. A. (1996). Studying context: A comparison of activity theory, situated action models, and distributed cognition. In B. A. Nardi (Ed.), *Context and consciousness: Activity theory and human–computer interaction* (pp. 69–102). Cambridge, MA: MIT Press.

Nelson, J. B., & Bouton, M. E. (2002). Extinction, inhibition, and emotional intelligence. In L. F. Barrett and P. Salovey (Eds.), *The wisdom in feeling: Psychological processes in emotional intelligence* (pp. 60–85). New York, London: The Guilford Press.

Nestler, E. J. (2001). Molecular basis of long-term plasticity underlying.*science, 2,* pp. 119–128.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman and S. W. Draper (Eds.), *User centered system design: New perspectives on human–computer interaction* (pp. 31–61). Hillsdale, NJ: Lawrence Erlbaum Associates.

Norman, D. A. (1988). *The design of everyday things.* New York: Doubleday Currency.

Norman, D. A. (1991). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human–computer interface* (pp. 17–38). Cambridge: Cambridge University Press.

Norman, D. A. (1993). *Things that make us smart.* Reading, MA: Addison-Wesley.

Norman, D (1999). Affordance, conventions and design. *Interactions, May/June,* pp. 38–42.

O'Donoghue, T., & Rabin, M. (2003). Self-awareness and self-control. In G. Loewenstein, D. Read and R. Baumester (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 217–243). New York: Russell Sage Foundation.

Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review, 108,* pp. 483–522.

O'Regan, J. K., & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24(5),* pp. 939–1011.

Paller, K. A., & Wagner, A. D. (2002). Observing the transformation of experience into memory. *Trends in Cognitive Science, 6(2),* pp. 93–102.

Parson, R., & Hartig, T. (2000). Environmental psychophysiology. In J. T. Cacioppo, L. G. Tassinary and G. G. Berntson (Eds.), *Handbook of psychophysiology* (pp. 815–846). Cambridge: Cambridge University Press.

Petroski, H. (1992). *The evolution of useful things.* New York: Vintage Books.

Petroski, H. (1999). *The book on the book shelf.* New York: Alfred A. Knopf, Inc.

Pfaffenberger, B. (1988). Fetished objects and humanized nature: Towards an anthropology of technology. *Man, New Series, 23(2),* pp. 236–252.

Pfaffenberger, B. (1992). Social anthropology of technology. *Annual Review of Anthropology, 21,* pp. 491–516.

Preece et al. (1994). *Human–computer interaction.* Wokingham, England: Addison-Wesley.

Quilici-Pacaud, J. -F. (1993). Dominant representations and technical choices: A method of analysis with examples from aeronautics. In P. Lemonnier (Ed.), *Technological choices* (pp. 399–412). London and New York: Routledge.

Rachlin, H. (2000). *The science of self-control.* Boston: Harvard University Press.

Read, D. (2003). Subadditive intertemporal choice. In G. Loewenstein, D. Read and R. Baumester (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 301–322). New York: Russell Sage Foundation.

Resnick, L. B. (1991). Shared cognition: Thinking as social practice. In L. B. Resnick, J. M. Levine and S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 1–20). Washington DC: American Psychological association.

Rhodes, B. J. (2000). Margin notes: Building a contextually aware associative memory. In *Proceedings of the 2000 international conference on intelligent user interfaces* (pp. 219–224). New Orleans, LA: ACM Press.

Ribeiro, G. (1996). Situating action III: Acting, dwelling, and squatting: An ecological approach to the relation between person and urban environment. *Ecological Psychology, 8(2),* pp. 131–151.

Robinson, T. E., & Berridge, K, C. (2003). Addiction. *Annual Review of Psychology, 54,* pp. 25–53.

Rogers, Y. & Ellis, J. (1994). Distributed cognition: an alternative framework for analysing and explaining collaborative working. *Journal of Information Technology, 9(2),* pp. 119–128.

Rolls, E. T. (1999). *The brain and emotion.* Oxford: Oxford University Press.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland and D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, vol. 2* (pp. 7–57). Cambridge, MA, MIT Press.

Sally, D. (2000a). Confronting the sirens: Rational behavior in the face of changing preferences. *Journal of Institutional and Theoretical Economics, 156(4)*, pp. 685–714.

Sally, D. (2000b). I, too, sail past: Odysseus and the logic of self-control. *Kyklos, 53(2),* pp. 173–200.

Salomon, G. (1993). No distribution without individuals' cognition: a dynamic interactional view. In G. Salomon (Ed.), *Distributed cognitions – Psychological and educational considerations* (pp. 111–138). Cambridge: Cambridge University Press.

Schacter, D. L. & Dodson, C. S. (2001). Misattribution, false recognition and the sins of memory. *Philosophical Transactions of the Royal Society, 356,* pp. 1385–1393.

Schiffer, M. B. (1992). *Technological perspectives on behavioral change.* Tucson: University of Arizona Press.

Scribner, S. (1986). Thinking in action: some characteristics of practical thought. In R. Sternberg and R. Wanger (Eds.), *Practical intelligence: Nature and origins of competence in the everyday world* (pp. 13–30). Cambridge: Cambridge University Press.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7(1–3),* pp 1–15.

Singletary, B., & Starner, T. (2001). Learning visual models of social engagement. In *Second international workshop on rec-*

ognition, analysis and tracking of faces and gestures in real-time systems, 2001, Vancouver, Canada.

Stammers, R. B., Carey, M. S., & Astley, J. A. (1990). Task analysis. In J. R. Wilson and E. N. Corlett (Eds.), *Evaluation of human work* (pp. 134–160). London: Taylor and Francis.

Staudenmaier, J. M. (1984). What SHOT hath wrought and what SHOT hath not: Reflections on twenty-five years of the history of technology. *Technology and Culture, 25(4),* pp. 707–730.

Staudenmaier, J. M. (1990). Recent trends in the history of technology. *American Historical Review, 95(3),* pp. 715–725.

Suchman, L. A. (1987). *Plans and situated actions: The problem of human–machine communication.* Cambridge: Cambridge University Press.

Swain, M. J., Kahn, R. E., & Ballard, D. H. (1992). Low resolution cues for guiding saccadic eye movements. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (pp. 737–740).

Tamada, J. A., Lesho, M., & Tierny, M. J. (2002). Keeping watch on glucose. *IEEE Spectrum Online,* April.

von Uexküll, J. (1934/1992). A stroll through the worlds of animals and men: A picture book of invisible worlds. Semiotica 89(4), pp. 319–391.

Walenstein, A. (2002). Foundations of cognitive support: Towards abstract patterns of usefulness. In P. Forbrig, Q. Limbourg, B. Urban and J. Vanderdonckt (Eds.), *Interactive systems design, specification, and verification, 9th international workshop, DSV-IS 2002* (pp. 133–147). Heidelberg: Springer-Verlag.

Warren, C. W. & McDonough, B. E. (1999). Event-related brain potentials as indicators of smoking cue-reactivity. *Clinical Neurophysiology, 110,* pp. 1570–1584.

Wertsch, J. V. (1998). *Mind as action.* Oxford: Oxford University Press.

WHO (2001). *The world health report. Mental health: New understanding, new hope.* Geneva: World Health Organization.

de Wit, O., van den Ende, J., Schot, J., & van Oost, E. (2002). Innovation junctions: Office technologies in the Netherlands, 1880–1980. *Technology and Culture, 43,* pp. 50–72.

Woods, D. D. & Roth, E. M. (1988). Cognitive systems engineering. In M. Helander (Ed.), *Handbook of human–computer interaction* (pp. 3–43). New York: North-Holland.

World Bank (1999). *Development in practice – Curbing the epidemic: governments and the economics of tobacco control.* Washington DC: World Bank.

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science, 21,* pp. 179–217.

Zhang, J. & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science, 18,* pp. 87–122.

Ziman, J. (Ed.), (2000). *Technological innovation as an evolutionary process*. Cambridge: Cambridge University Press.