To appear in ARIST: Annual Review of Information Science and Technology, no 38, 2004

NEW THEORETICAL APPROACHES FOR HCI

Yvonne Rogers
Interact Lab, School of Cognitive and Computing Sciences, University of Sussex
Brighton, BN1 9QH, UK; email: yvonner@cogs.susx.ac.uk

Theory weary, theory leery, why can't I be theory cheery? (Erickson, 2002, p269)

The field of human-computer interaction is rapidly expanding. Alongside the extensive technological developments that are currently taking place, is the emergence of a 'cottage industry' culture, where a polyphony of new theories, methods and concerns have been imported into the field from a diversity of disciplines and backgrounds. An extensive critique of recent theoretical developments is presented together with what practitioner's currently use. A significant development of importing new theories into the field has been much insightful explication of 'HCI' phenomena, together with extending the field's discourse. However, at the same time, the theoretically-based approaches have had a limited impact on the practice of interaction design. This chapter discusses why this is so and suggests that different kinds of mechanisms are needed that will enable both designers and researchers to better articulate and theoretically ground the hard challenges facing them today.

Introduction

The field of human-computer interaction is bursting at the seams. Its mission, raison d'être, goals and methodologies, that were well established in the 80s, have all greatly expanded to the point that "HCI is now effectively a boundless domain" (Barnard *et al.,* 2000, p221). Everything is in a state of flux: the theory driving the research is changing, a flurry of new concepts are emerging, the domains and type of users being studied are diversifying, many of the ways of doing design are new and much of what is being designed is significantly different. While potentially much is to be gained from such rapid growth, the downside is an increasing lack of direction, structure and purpose in the field. What was originally a confined problem space with a clear focus that adopted a small set of methods to tackle it – that of designing computer systems to make them more easy and efficient to use by a single user – is now turning into a more diffuse problem space with a less clear purpose as to what to study, what to design for and which methods to use. Instead, aspirations of overcoming the Digital Divide, through providing universality and accessibility for all, have become driving concerns (e.g. Shneiderman, 2002a). It comes as no surprise that the move towards more openness is, likewise, happening in the field, itself. Many more topics, areas and approaches are now considered acceptable research and practice.

A problem with allowing a field to expand in this eclectic way is that it can easily get out of control. No-one really knows what its purpose is anymore or indeed what criteria to use to assess its contribution and value to knowledge and practice. For example, of all the many new approaches, ideas, methods and goals that are now being proposed how do we know which are acceptable, reliable, useful and generalisable? Moreover, how do researchers and designers, alike, know which of the many tools and techniques to use when doing design and research? What do they use to help make such judgments?

To be able to address these concerns, a young field in a state of flux (as is HCI) needs to take stock and begin to reflect on the numerous changes that are happening. The purpose of this chapter is to consider theoretical developments, by assessing and reflecting upon the role of theory in contemporary HCI and the extent to which it is used in design practice. Over the last ten years, a diversity of new theories have been imported and adapted into the field. A key question raised is whether such attempts have been productive in terms of 'knowledge transfer'. By knowledge transfer, it is meant here the translation of research findings (e.g. theory, empirical results, descriptive accounts, cognitive models) from one discipline (e.g. cognitive psychology, sociology) into practical concerns that can be applied to another (e.g. Human-Computer Interaction, Computer Supported Cooperative Work).

Why the explosive growth in HCI?

One of the main reasons for the dramatic change in direction in HCI is as a reaction to the explosion of new challenges confronting it. The arrival and rapid pace of technological developments in the last few years (e.g. the internet, wireless technologies, handheld computers, wearables, pervasive technologies, tracking devices) has led to an escalation of new opportunities for augmenting, extending and supporting user experiences, interactions and communications. These include designing experiences for all manner of people (and not just *users*) in all manner of settings doing all manner of things. The home, the crèche, the outdoors, public places and even the human body are now being experimented with as potential places to embed computational devices. Furthermore, a far reaching range of human activities is now being analyzed and technologies proposed to support them, even to the extent of invading previously private and taboo aspects of our lives (e.g. domestic life and personal hygiene). A consequence is that 'the interface' is becoming ubiquitous. Computer-based interactions can take place through many kinds of surfaces and in many different places. As such, many radically different ways of interacting with computationally-based systems are now possible, ranging from the visible that we are conscious of (e.g. using a keyboard with a computer monitor) to the invisible that we are unaware of (e.g. our physical movements triggering toilets to flush automatically through sensor technology).

In an attempt to keep up and appropriately deal with the new demands and challenges, significant strides have been made in academe and industry, alike, towards developing an armory of methodologies and practices. Innovative design methods, unheard of in the 80s, have been imported and adapted from far a field to study and investigate what people do in diverse settings. Ethnography, informant design, cultural probes and scenario-based design are examples of these (see Rogers *et al.,* 2002). New ways of conceptualizing the field are also emerging. For example, usability is being operationalized quite differently, in terms of a range of user experience goals (e.g. aesthetically pleasing, motivating, fun) in addition to the traditional set of efficiency goals (Rogers, *et al.,* op cit). The name interaction design is also increasingly being banded about in addition to of human-computer interaction, as a way of focusing more on what is being done (i.e. designing interactions) rather than the components it is being done to (i.e. the computer, the human). This more encompassing term generally refers to:

"the design of interactive products to support people in their everyday and working lives" (Rogers, *et al.,* 2002, p.6) and "designing spaces for human communication and interaction" (Winograd, 1997 p. 155).

New paradigms for guiding interaction design are also emerging. The prevailing desktop paradigm, with its concomitant GUI and WIMP interfaces, is being superseded by a range of new paradigms, notably ubiquitous computing ('UbiComp'), pervasive environments and everyday computing. The main thrust behind the paradigm of ubiquitous computing came from the late Mark Weiser (1991), whose vision was for computers to disappear into the environment in a way that we would no longer be aware of them and would use them without thinking about them. Similarly, a main idea behind the pervasive environments approach is that people should be able to access and interact with information any place and any time using a seamless integration of technologies.

Alongside these methodological and conceptual developments, has been a major rethink of whether, how and what kinds of theory can be of value in contributing to the design of new technologies. On the one hand, are strong advocates, arguing that there definitely needs to be a theoretical foundation to address the difficult design challenges ahead that face the HCI community (e.g. Barnard *et al.,* 2000; Hollan *et al.,* 2000; Kaptelinin, 1996; Sutcliffe, 2000) and that, furthermore, there is a distinct lack of it currently in the field (Castel, 2002). On the other, there are those who argue that theory has never been useful for the practical concerns of HCI and that it should be abandoned in favor of continuing to develop more empirically-based methods to deal with the uncertain demands of designing quite different user experiences using innovative technologies (e.g. Landauer, 1991). In this chapter, I examine the extent to which early and more recent theoretical developments in HCI have been useful and then contrast this with two surveys that examine the extent to which they have been useful in the practice of doing interaction design.

Early theoretical developments in HCI

In the early '80s, there was much optimism as to how the field of cognitive psychology could significantly contribute to the development of the field of HCI. A driving force was the realization that most computer systems being developed at the time were difficult to learn, difficult to use and did not enable the users to carry out the tasks in the way they wanted. The body of knowledge, research findings and methods that made up cognitive psychology were seen as providing the means by which to reverse this trend, by being able to inform the design of easy to learn and use computer systems. Much research was carried out to achieve this goal: mainstream information processing theories and models were used as a basis from which to develop design principles, methods, analytic tools and prescriptive advice for the design of computer interfaces (e.g. see Carroll, 1991). These can be loosely classified into three main approaches: applying basic research, cognitive modeling and the populist dissemination of knowledge.

Applying basic research: Early attempts at using cognitive theory in HCI brought in relevant theories and appropriated them to interface design concerns. For example, theories about human memory were used to decide what were the best set of icons or command names to use, given people's memory limitations. One of the main benefits of this approach was to help researchers identify relevant cognitive factors (e.g. categorization strategies, learning methods, perceptual processes) that are important to consider in the design and evaluation of different kinds of GUIs and speech recognition systems.

A core lesson that was learned, however, is that you cannot simply lift theories out of an established field (i.e. cognitive psychology), that have been developed to explain specific phenomena about cognition, and then reapply them to explain other kinds of seemingly related phenomena in a different domain (i.e. interacting with computers). This is because the kinds of cognitive processes that are studied in basic research are quite different from what happens in the 'real' world of human-computer interactions (Landauer, 1991). In basic research settings, behavior is controlled in a laboratory in an attempt to determine the effects of singled out cognitive processes (e.g. short term memory span). The processes are studied in isolation and subjects (sic) are asked to perform a specific task, without any distractions or aids at hand. In contrast, the cognition that happens during human-computer interaction is much more 'messy', whereby many interdependent processes are involved for any given activity. Moreover, in their everyday and work settings, people rarely perform a task in isolation. Instead, they are constantly interrupted or interrupt their own activities, by talking to others, taking breaks, starting new activities, resuming others, and so on. The stark differences between a controlled lab setting and the messy real world setting, meant that many of the theories derived from the former were not applicable to the latter. Predictions based on basic cognitive theories about what kinds of interfaces would be easiest to learn, most memorable, easiest to recognize and so on, were often not supported.

The problem of applying basic research in a real world context is exemplified by the early efforts of a number of cognitive psychologists in the early 80s, who were interested in finding out what was the most effective set of command names for text editing systems, in terms of being easy to learn and remember. At the time, it was a well-known problem that many users and some programmers had a difficult time remembering the names used in command sets for text editing applications. Several psychologists assumed that research findings on paired-associate learning could be usefully applied to help overcome this problem; this being a well developed area in the basic psychological literature. One of the main findings that was applied was that pairs of words are learned more quickly and remembered if subjects have prior knowledge of them (i.e. highly familiar and salient words). It was further suggested that command names be designed to include specific names that have some natural link with the underlying referents they were to be associated with. Based on these hypotheses, a number of experiments were carried out, where users had to learn different sets of command names, that were selected based on their specificity, familiarity, etc. The findings from the studies, however, were inconclusive; some found specific names were better remembered than general terms (Barnard *et al.,* 1982), others showed names selected by users, themselves, were preferable (e.g. Ledgard *et al.,* 1981; Scapin, 1981) while others demonstrated that high frequency words were better remembered than low frequency ones (Gunther *et al.,* 1986). Hence, instead of the outcome of the research on command names being able to provide a generalisable design rule about which names are the most effective to learn and remember, it suggested that a whole range of different factors affects the learnability and memorability of command names. As such, the original theory about naming was not able be applied effectively to the selection of optimal names in the context of computer interfaces.

Cognitive modeling: Another attempt to apply cognitive theory to HCI, was to model the cognition that is assumed to happen when a user carries out their tasks. Some of the earliest models focused on user's goals and how they could achieve (or not) them with a particular computational system. Most influential at the time were Hutchins *et al.'s* (1986) conceptual framework of directness, which describes the gap between the user's goals and the way a system works in terms of gulfs of execution and evaluation, and Norman's (1986) theory of action, which models the putative mental and physical stages involved in carrying out an action when using a system. Both were heavily influenced by contemporary cognitive science theory of the time, which itself, focused on modeling people's goals and how they were met.

The two cognitive models essentially provided a means by which to conceptualize and understand the interactions that were assumed to take place between a user and a system. In contrast, Card *et al's* (1983) model of the user, called the model human processor (MHP), went further by providing a basis from which to make quantitative predictions about user performance and, in so doing, provided a means by which to allow researchers and developers to evaluate different kinds of interfaces to assess their suitability for supporting various tasks. Based upon the established information processing model of the time (that, itself, had been imported into cognitive psychology), the MHP comprised interacting perceptual, cognitive and motor systems, each with their own memory and processor. To show how the model could be used to evaluate interactive systems, Card *et al.* developed a further set of predictive models, collectively referred to as GOMS (Goals, Operators, Methods and Selection rules).

Since its inception, a number of researchers have used and extended GOMS, reporting on its success for comparing the efficacy of different computer-based systems (see Olson and Olson, 1991). Most of these have been done in the lab but there have been a few carried out in a real-world context. The most well-known is Project Ernestine, where a group of researchers carried out a GOMS analysis for a modern workstation that a large phone company were contemplating purchasing, and counter-intuitively, predicted that it would perform worse than the existing computer system being used at the company, for the same kind of tasks. A consequence was that they advised the company not to invest in what could have been potentially a very costly and inefficient technology (Attwood *et al.,* 1996). While this study has shown that the GOMS approach can be useful in helping make decisions about the effectiveness of new products, it is not often used for evaluation purposes (although there is some evidence of wider use in the military). Part of the problem is its highly limited scope: it can only reliably model computer-based tasks that involve a small set of highly routine data-entry type tasks. Furthermore, it is intended to be used to predict expert performance, and does not allow for errors to be modeled. This makes it much more difficult (and sometimes impossible) to predict how most users will carry out their tasks when using systems in their work, especially those that have been designed to be flexible in the way they can be used. In most situations the majority of users are highly variable in how they use systems, often carrying out their activities in quite different ways to that modeled or predicted. Many *unpredictable* factors come into play. These include individual differences among users, fatigue, mental workload, learning effects and social and organizational factors (Olsen and Olsen, 1991). Moreover, most people do not carry out their tasks sequentially but tend to be constantly multi-tasking, dealing with interruptions and talking to others, while carrying out a range of activities. A problem with using predictive models, therefore, is that they can only make predictions about isolated predictable behavior. Given that most people are often unpredictable in the way they behave and, moreover, interweave their ongoing activities in response to unpredictable external demands, it means that the outcome of a GOMS analysis can only ever be a rough approximation and sometimes even be inaccurate. Furthermore, many would argue that carrying out a simple user test, like heuristic evaluation, can be a more effective approach that takes much less effort to use (see table 1).

Table 1 Time it takes to train and effort involved for different analytic methods in HCI (adapted from Olson and Moran, 1996, p.281)

Method	Effort	Training	Reference
Checklists (e.g. heuristic evaluation)	1 day	1 week	Shneiderman,1992
Cognitive walkthrough	1 day	3 months	Lewis *et al.,* 1990
Cognitive complexity theory	3 days	1 year	Kieras, 1988
GOMS	3 days	1 year	Card *et al.,* 1983

Despite the disparity between the outcome of a modeling exercise and the vagaries of everyday life, a number of other cognitive models have been developed, aimed at predicting user behavior when using various kinds of systems (e.g. the EPIC model, Kieras and Meyer, 1997). Similar to the various versions of GOMS, they can predict simple kinds of user interaction fairly accurately, but are unable to cope with more complex situations, where the amount of judgment a researcher or designer has to make, as to which aspects to model and how to do this, greatly increases (Sutcliffe, 2000). The process becomes increasingly subjective and involves considerable effort, making it more difficult to use them to make predictions that match the ongoing state of affairs.

In contrast, cognitive modeling approaches, that do not have a predictive element to them, have proven to be more successful in their utility in practice. Examples include heuristic evaluation (Mohlich and Nielsen, 1990) and cognitive walkthroughs (Polson *et al,* 1992) which are much more widely used by practitioners. Such methods provide various heuristics and questions for evaluators to operationalize and answer, respectively. An example of a well known heuristic is 'minimize user memory load'. As such, these more pragmatic methods differ from the other kinds of cognitive modeling techniques insofar as they provide *prescriptive* advice, that is largely based on assumptions about the kinds of cognitive activities users engage in when interacting with a given system. Furthermore, their link to a theoretical basis is much looser.

Diffusion of popular concepts: Perhaps, the most significant and widely-known contribution that the field of cognitive psychology made to HCI is the provision of *explanations* of the capabilities and limitations of users, in terms of what they can and cannot do when performing computer-based tasks. For example, theories that were developed to address key areas, like memory, attention, perception, learning, mental models and decision-making have been much popularized in tutorials, introductory chapters, articles in magazines and the web, to show their relevance to HCI. Examples of this approach include Preece *et al.* (1994), Norman (1988) and Monk (1984). By explicating user performance in terms of well known cognitive characteristics that are easy to assimilate (e.g. recognition is better than recall), designers can be alerted to their possible effects when making design decisions – something that they might not have otherwise considered. A well known example is the application of the finding that people find it easier to recognize things shown to them than to have to recall them from memory. Most graphical interfaces have been designed to provide visual ways of presenting information, that enable the user to scan and recognize an item like a command, rather than require them to recall what command to issue next at the interface.

This approach, however, has tended to be piecemeal – depending on the availability of research findings in cognitive psychology that can be translated into a digestible form. A further problem with this approach is its propensity towards a 'jewel in the mud' culture, whereby a single research finding sticks out from the others and is much cited, at the expense of all the other results (Green *et al.,* 1996). In HCI, we can see how the 'magical number 7+-2' (George Miller's theory about memory, which is that only 7+-2 chunks of information, such as words or numbers, can ever be held in short term memory at any one time) has become the de facto example: nearly every designer has heard of it but not necessarily where it has come from or what situations it is appropriate to apply. A consequence is that it has largely devolved into a kind of catch-phrase, open to interpretation in all sorts of ways, which can end up being far removed from the original idea underlying the research finding. For example, some designers have interpreted the magic number 7+-2 to mean that displays should have no more than 7+-2 of a category (e.g. number of colors, number of icons on a menu bar, number of tabs at the top of a web page and number of bullets in list), regardless of context or task, which is clearly in many cases inappropriate (see Bailey, 2000).

A shift in thinking

We have examined the ways in which cognitive theory was first applied in HCI. These can be classified, largely, as:

- informative (providing useful research findings)
- predictive (providing tools to model user behavior)
- prescriptive (providing advice as to how to design or evaluate)

In the late 80s, however, it became increasingly apparent that these early attempts were limited in their success; neither matching nor scaling up to the demands and perceived needs of developing systems. Several researchers began to reflect on why the existing theories, that had been imported from cognitive psychology, were failing to be more widely applied to the problems of design and computer use (e.g. Long and Dowell, 1996). Much criticism was expressed about the inadequacies of classical cognitive theories for informing system design (e.g. see Carroll, 1991). A number of problems were identified, including that the theories were too low-level, restricted in their scope and failed to deal with real world contexts (Barnard, 1991). There was much concern, leading to calls to abandon what has been coined as the 'one-stream' approach, whereby it was naively assumed that mainstream theory provided by pure science (i.e. cognitive psychology) could trickle down into the applied science of designing computer systems (see Long and Dowell, 1996). There was even criticism that psychologists were merely using the field of HCI as a test bed for trying out their general cognitive theories (Bannon and Bødker, 1991) or for validating the assumptions behind specific models (Barnard and May, 1999). Instead, it was argued that other kinds of theories were needed that were more encompassing, addressing more directly the concerns of interacting with computers in real-world contexts. It was still assumed that theory *did* have a valuable role to play in helping to conceptualize the field, provided it was the right theory. The question was what kind of theory and what role should it play? By changing and dissolving the boundaries of what was under scrutiny, and by reconceptualizing the phenomena of interest, using different theoretical lenses and methods, it was further assumed that the pertinent issues in the field could be recast and in so doing, lead to the design of more usable computer artifacts (Bannon and Bødker, 1991).

Several researchers began searching elsewhere, exploring other disciplines for theories that could achieve this. An early contender that was put forward was Activity Theory originating from Soviet psychology (Bødker, 1989; Kuutti, 1996; Engestrøm and Middleton, 1996; Nardi, 1996). It was regarded as a unifying theoretical framework for HCI, being able to both provide the rigor of the scientific method of traditional cognitive science while taking into account social and contextual aspects (Kaptelinin et al, 1999). There were also attempts to look for theories that took into account how the environment affected human action and perception. Several ideas from ecological psychology were reconceptualized for use in the field (e.g. Gaver, 1991; Norman, 1988). At the same time, several researchers sought substantially to revise or adapt existing cognitive frameworks so as to be more representative and build directly on the concerns of HCI (e.g. Draper, 1993). Long and Dowell (1989, 1996) made persistent calls for more domain-specific theories that focus on the concerns of users interacting with computers to enable them to work effectively. Carroll *et al.* (1991) also advocated the need for this change in their task-artifact cycle framework, arguing that users and designers would benefit more if the process by which tasks and artifacts co-evolved could be "better understood, articulated and critiqued" (p.99). Two main approaches that have emerged from cognitive science are distributed cognition and external cognition. A central focus of these approaches is the structural and functional role of external representations and artifacts in relation to how they are used in conjunction with internal representations (e.g. *Green et al.,* 1996; Hutchins, 1995; Kirsh, 1997; Scaife and Rogers, 1996; Wright *et al.,* 2000).

There was also a 'turn to the social' (Button 1993): sociologists, anthropologists and others in the social sciences came into HCI, bringing new frameworks, theories and ideas about technology use and system design. These were primarily the situated action approach and ethnography. Human computer interactions was conceptualized as social phenomena (e.g. Heath and Luff, 1991). A main thrust of this approach was to examine the *context* in which users interact with technologies: or put in social terms, how people use their particular circumstances to achieve intelligent action. The approach known as ethnomethodology (Garfinkel, 1967; Garfinkel and Sacks, 1970), that had itself come about as a reaction against mainstream sociology, provided much of the theoretical and methodological underpinning (Button, 1993). In particular, it was assumed that ethnomethodology could offer descriptive accounts of the informal aspects of work (i.e. "the hurly burly of social relations in the workplace and locally specific skills required to perform any task", Anderson, 1994, p154) to *complement* the formal methods and models of software engineering and in so doing, begin to address some of the 'messiness' of human technology design mentioned at the beginning of the chapter, and which cognitive theories have not been able to adequately address.

How have recent theoretical approaches fared in the field?

In this section, I examine in more detail how recent theoretical developments in HCI have fared. In particular, I look at how researchers have attempted to transform alternative theoretical knowledge into an applied form – aimed at being used by others – especially practitioners. Concomitantly, I look at whether and how people, who develop and evaluate technologies and software (e.g. designers, usability practitioners, information architects), have used them. In particular, I consider whether researchers have been successful in providing a new body of theoretical knowledge that is tractable to others. I begin by looking at the most well received and referenced attempts at importing different kinds of theory into HCI. I follow this by examining what a cross-section of practitioners use in their work and which of the new approaches they have found useful.

The researcher's perspective

Below I analyze the contributions made by researchers in HCI for the following: ecological approach, activity theory, external cognition, distributed cognition, situated action, ethnomethodology[1], hybrid and overarching approaches. The reason for the selection of these approaches is that they are considered to be the main ones that have been imported, applied and developed in HCI over the last 10-15 years. As such, it is not meant to be an exhaustive list of developments in the field, but an attempt to show how the recent generation of theories and approaches have been developed, transformed, and applied to practical concerns.

The ecological approach

The ecological approach evolved primarily from Gibson's (1966, 1979) view that psychology should be the study of the interaction between the human and its environment. Its concern is with providing a carefully detailed description of the environment and people's ordinary activities within it (Neisser, 1985). A number of researchers within HCI have adapted the approach for the purpose of examining how people interact with artifacts. These include Gaver (1991), Kirsh, (2001), Norman (1988), Rasmussen and Rouse (1981), Vicente (1995) and Woods (1995).

A main focus in the original ecological framework, was to analyze invariant structures in the environment in relation to human perception and action. From this framework two key related concepts have been imported into HCI: *ecological constraints* and *affordances*. Of the two, the latter is by far the most well known in HCI. Ecological constraints refer to structures in the external world that guide people's actions rather than those that are determined by internal cognitive processes. The term affordances, within the context of HCI, has been used to refer to attributes of objects that allow people to know how to use them. In a nutshell, to afford is taken to mean 'to give a clue' (Norman, 1988). Specifically, when the affordances of an object are perceptually obvious it is assumed that they make it easy to know how to interact with the object (e.g. door handles afford pulling, cup handles afford grasping). Norman (1988) provides a range of examples of affordances associated with everyday objects such as doors and switches.

This explication of the concept of affordances is much simpler than Gibson's original idea. One of the main differences is that it only refers to the properties of an object, whereas Gibson used it to account for the relationship between the properties of a person and the perceptual properties of an object in the environment. An obvious advantage of simplifying it in this manner, is that it makes it more accessible to those not familiar with Gibsonian ideas. Indeed, this way of thinking about affordances has been much popularized in HCI, providing a way of describing properties about interface objects that highlight the importance of making 'what can be done to them' obvious. One suggestion is that this reformulation helps designers think about how to represent objects at the interface that will readily afford permissible actions (Gaver, 1991) and provide cues as to how to interact with interface objects more easily and efficiently. However, a problem of appropriating the concept of affordance in this manner is it puts the onus on the designer

---

[1] It should be noted that ethnomethodology is viewed as atheoretical, and has been imported into HCI primarily as an analytic approach

to use their intuition as to how to decide what are affordable objects at the interface (St Amant, 1999). There are no abstractions, methods, rules or guidelines to help them – only analogies drawn from the real world. The lack of guidance has unfortunately led to the concept being somewhat glibly used:

> "I put an affordance there," a participant would say, "I wonder if the object affords clicking…" affordances this affordances that. And no data, just opinion, Yikes! What had I unleashed upon the world?"
>
> Don Norman's (1999, p38) reaction to a recent CHI-Web discussion.

Furthermore, in its borrowed form, the concept of affordance has often been interpreted in a design context as suggesting that one should try to emulate real-world objects at the interface – which is clearly a far cry from Gibson's ideas and is highly questionable. The increasing trend towards bringing high fidelity realism to the interface (i.e. designing objects to appear as 3D at the interface to give the illusion of behaving and looking like real world counterparts) is witness to this. On-screen buttons are increasingly being designed now to have a 3D look, to give the appearance of protruding. An assumption is that this kind of representation will give the buttons the affordance of pushing, inviting the user to click on them, in an analogous way to what they would do with actual physical buttons. While users may readily learn this association, it is equally the case that they will be able to learn how to interact with a simple, 2D representation of a button on the screen. The effort to learn the association is likely to be similar. In addition, it is not always the case that 3D buttons are the most effective form of representation. For example, simple, plain and abstract representations may prove to be far easier to recognize and distinguish from each other for applications where there are many operations and functions that need to be represented at the interface (e.g. CAD).

Norman (1999) has since tried to deal with the pervasive misunderstanding and misuse of the term, since his original explication of it in his POET book (Norman, 1988). In its place, he now argues for two kinds of affordance: perceived and real. Physical objects are said to have real affordances, as described above, like grasping, which are perceptually obvious and do not have to be learned. In contrast, user-interfaces, that are screen based, do not have these kinds of real affordances, meaning that the user needs to learn the meaning and function of each object represented at the interface before knowing how to act. Therefore, it does not make sense to talk about interface design in terms of real affordances. Alternatively, Norman argues that screen-based interfaces have perceived affordances, which are based on learned conventions and feedback. For example, having a red flashing button icon appear at the interface may provide visual cues to enable the user to perceive that clicking on that icon is a meaningful useful action at that given time in their interaction with the system, that has a known outcome. However, this begs the question of where the 'ecology' has gone, namely the unconscious cue-action coupling that underlies the true sense of the term affordance.

The downside of the concept of affordance being popularized in this way is that the richness and contextual background of the original theory has been lost, making it difficult to appreciate its significance other than at a superficial level. Some may argue that this does not matter since it has provided designers with a new way of thinking and talking about design that they did not have before. However, others would argue that it can distort their way of thinking about interaction design to the extent that it overly constrains the way they do design, as satirized by Norman in his CHI-website quote.

One way of putting the currency back into the concept may be to try to import more knowledge about what is meant by it. Kirsh (2001), for example, describes the notion of affordance in terms of entry points, which refer to the way structures in the environment *invite* people to do something. For example, the way information is laid out on posters, websites and magazines provides various entry points for scanning, reading and following. These include headlines, columns, pictures, cartoons, figures, tables, icons, etc. Well designed information allows a person's attention to move rapidly from entry point to entry point for different sections (e.g. menu options, lists, descriptions). In contrast, poorly designed information does not have clear entry points – it is hard to find things. In Kirsh's terms, entry points are like affordances, inviting people to carry out an activity (e.g. read it, scan it, look at it, listen to it, click on it). This reconceptualization potentially has more utility as a design concept insofar as it gives more clues as to what to do with it: encouraging designers to think about the coordination and sequencing of actions and the kind of feedback to provide, in relation to how objects are positioned and structured at an interface – rather than simply whether objects *per se* afford what to do with them.

Another attempt at pulling in more of the original theory has been to develop extensive frameworks, focusing more on the notion of ecology and what it means for design. For example, Vicente (1995) and Vicente and Rassmussen (1990) have developed the Ecological Interface Design framework (EID), where they describe affordances in terms of a number of actions (e.g. moving, cutting, throwing, carrying). The various actions are sorted into a hierarchy of categories, based on what, why and how they afford. The outcome is a framework which is intended to allow designers to analyze a system at different levels, which correspond to the levels in the hierarchy. St Amant (1999) has also attempted to develop an ecological framework, where he specifies a number of different kinds of affordances in relation to planning representations, derived from AI research. He suggests that his framework can "contribute to an understanding of low level actions in a graphical user interface" (p333). However, it is not clear how much of the two frameworks is ecologically-based. In both, there is much more emphasis on modeling user's actions *per se* rather than the ecological interactions between a person and their environment. As such, the sense of perceptual coupling is lost. Moreover, other cognitive theoretical frameworks, like Rasmussen's (1986), seems to play a much greater contribution. Although the frameworks may prove to be useful tools, it cannot be said to be due to any theoretical insights gained from Ecological Psychology.

To summarize, a main contribution of the ecological approach for HCI has been to extend its discourse, primarily in terms of articulating certain properties about objects at the interface in terms of their behavior and appearance. As such the role of theory, here, is largely *descriptive,* providing a key design concept. The affordance (sic) of the term affordance has led to it becoming one of the most common terms used in design parlance. Less familiar and, so far, less used is the theory as an *analytic framework*, by which to model human activities and interactions. In the next section, I discuss how the Activity Theory approach has been developed as an analytic framework and examine how useful it has been.

The Activity Theory approach

Activity theory has its origins in Soviet Psychology (Leontiev, 1978). Its conceptual framework was assumed to have much to offer to HCI, in terms of providing a means of analyzing actions and interactions with artifacts within a historical and cultural context – something distinctly lacking in the cognitive paradigm (Bannon and Bødker, 1991; Bødker, 1989; Kuutti, 1996; Nardi, 1996). There are several introductions to the approach showing its potential relevance to HCI (e.g. Bannon and Bødker, 1991; Kaptelinin and Nardi, 1997; Kuutti, 1996) and a corpus of studies that have used its framework to analyze different work settings and artifacts-in-use. These include studies of user-interfaces for systems to be used in newspaper production (Bødker, 1989) and medical care in hospitals (Engestrøm, 1993) together with shaping the design of educational technology (Bellamy, 1996) and groupware (Fjeld *et al.,* 2002).

The purpose of Activity Theory in its original Soviet context was to explain cultural practices (e.g. work, school) in the developmental, cultural and historical context in which they occur, by describing them in terms of 'activities'. The backbone of the theory is presented as a hierarchical model of activity which frames consciousness at different levels, in terms of operations, actions and activities, together with a number of principles. A main rationale for bringing this particular framework into HCI was that it was considered useful for thinking about the design of user-interfaces and computer systems based in the work settings in which they were to be used (Bødker, 1989). It was also assumed that the theory could provide the contextual background that would allow technology to be designed and implemented that better suited workers in their work environments.

Since Bødker's initial application of the imported form of the theory, it has been used for a range of purposes in HCI, notably Kuutti (1996) extension of the hierarchical framework to show how information technology can be used to support different kinds of activities at different levels. Nardi (1996) has also used the framework to show how it can be of value for examining data and eliciting new sets of design concerns. Specifically, she recast data from a field study that she had carried out earlier to compare the benefits of task-specific versus generic application software for making slides (Nardi and Johnson, 1994). In doing this exercise second time round, but with the added benefit of the conceptual framework of activity theory at hand, she claimed to have been able to make more sense of her data. In particular, it enabled her to ask a more appropriate set of questions that allowed her subsequently to come up with an alternative set of recommendations about software architectures for the application of slide-making.

The most cited application of activity theory of recent is Engestrøm's (1990) extension of it within the context of his particular field of research known as 'developmental work research'. His framework was designed to include other concepts (e.g. contradictions, community, rules and division of labor) that were pertinent to work contexts and which could provide conceptual leverage for exploring these. Using this extended form of the framework, called the Activity System Model (see figure 1), he and his colleagues have analyzed a range of work settings – usually where there is a problem with existing or newly implemented technology – providing both macro and micro level accounts. Several others have followed Engestrøm's example and have used the model to identify a range of problems and tensions in various settings. Some have taken this variant and adapted it further to suit their needs. These include Halloran *et al.'s* (2002) Activity Space framework for analyzing collaborative learning, Spasser's (2002)'realist' approach for analyzing the design and use of digital libraries and Collins *et al's* (2002) model employed to help identify user requirements for customer support engineers. One of the putative benefits from having a more extensive framework with a set of conceptual foci is how they structure and scaffold the researcher/designer in their analysis:

"We found that activity system tensions provide rich insights into system dynamics and opportunities for the evolution of the system." (Collins *et al.,* op cit, p.58).

Figure 1 (i) The basic Activity Theory Framework and (ii) Engestrøm's (1987) extended Activity System Model

In many ways, the extended framework has proven attractive because it offers a "rhetorical force of naming" (Halverson, 2002, p247); providing an armory of terms that the analyst can use to match to instances in their data and, in so doing, systematically identify problems. However, such an approach relies largely on the analyst's interpretative skills and orientation as to what course to take through the data and how to relate this to which concepts of the framework. In many ways this is redolent of the problem discussed earlier concerning the application of cognitive modeling approaches to real world problems. There is little guidance (since it essentially is a subjective judgment) to determine the different kinds of activities – a lot depends on understanding the context in which they occur. It is argued, therefore, that to achieve a level of competence in understanding and applying activity theory requires considerable learning and experience. Hence, while, the adapted version of the activity system model and its variants have proven to be useful heuristic tools, they are really only useful for those who have the time and ability to study activity theory in its historic context. When given to others not familiar with the original theory, its utility is less productive. For example, the basic abstractions of the model, like object and subject, were found to be difficult to follow, and easily confused with everyday uses of the terms when used by design and engineering teams (who were initially unfamiliar with them) to discuss user requirements (Collins *et al.,* 2002).

In sum, the main role played by theory for this approach is *analytic,* providing a set of interconnected concepts that can be used to identify and explore interesting problems in field data.

The external cognition approach

As mentioned previously, one of the main arguments put forward as to why basic cognitive theories failed to make a substantial contribution to HCI was the mismatch between the cognitive framework (information processing model) and the phenomena of interest (i.e. human-computer interaction). The former had been developed to explain human cognition in terms of hypothetical processes exclusively inside the mind of one person. The latter is essentially about how people interact with external representations at the computer interface. As emphasized by Zhang and Norman (1994) "it is the interwoven processing of internal and external information that generates much of a person's intelligence" (p. 87). It is this interplay between internal and external representations that is the focus of the external cognition approach (Scaife and Rogers, 1996; see also Card *et al.,* 1999). An underlying aim has been to develop theoretical constructs that unite 'knowledge in the head' with 'knowledge in the world' (Norman, 1988; Vera and Simon, 1993; Wright *et al.,* 2000). In giving external representations a more central and functional role in relation to internal cognitive mechanisms, it is assumed that more adequate theoretical accounts of cognition can be developed.

A number of analytic frameworks have been developed that can be considered as part of the external cognition approach, and, in turn, various concepts have been operationalized to inform the design and evaluation of interactive technologies. For example, Green *et al.* (1996) developed a more complex model of cognitive processing by augmenting the original information processing one to take into account the dynamic interplay between inputs, outputs and processing. Zhang and Norman (1994) developed a theoretical framework of distributed representations for analyzing problem-solving behavior, where different combinations of external and internal representations are modeled in an abstract task space.

Similarly, Wright *et al.* (2000) modeled external cognition in terms of the putative abstract information types that are used and in so doing provided a set of interlinked theoretical constructs. These are labeled as 'resources' and categorized as being either plans, goals, possibilities, history, actions-effect relations or states. They can be represented internally (e.g. memorized procedure) or externally (e.g. written instructions). Configurations of these resources, distributed across internal and external representations, are assumed to be what informs an action. In addition, the way the resources are configured in the first place, is assumed to come about through various 'interaction strategies'. These include things like plan following and goal matching. Thus a user's selection of a given action may arise through an internal goal matching strategy (e.g. delete the file) being activated in conjunction with an external 'cause-effect relation' being perceived, (e.g. a dialog box popping up on the screen saying 'are you sure you want to delete this file?').

The thrust of Wright *et al's* (2000) cognitive model is to provide an analytic framework that can be used to determine the kinds of interaction that take place when a user interacts with a computer application. In some ways, it can be seen to have a rhetorical force that have parallels to the adapted frameworks of Activity Theory. Namely, there are several named concepts, that are linked through a relatively simple syntax, that allow observational data to be matched and modeled in them. In particular, the analyst can use the concepts to identify patterns and the variability of resources that are used at different stages of a task – such as determining when a user can depend on the external resources (e.g. action-effect relations) to constrain what to do next and when they must rely more on their own internal resources (e.g. plans, goals and history of actions). From this, the analyst can reflect on the problems with a given interface, in terms of the demands the various patterns of resources place on the user. In this sense, it is more akin to a traditional modeling tool, such as the cognitive task analytic methods discussed at the beginning of the chapter.

A different approach to applying theory arising from the external cognition approach is to provide a set of independent concepts that attempt to map a *theoretical space* specifically in terms of a *design space*. A number of design-oriented concepts have resulted, most notable, is the design vocabulary developed by Green (1989), called *cognitive dimensions,* that was intended to allow psychologists and importantly, others, to make sense of and use to talk together about design issues. Green's overarching goal was to develop a set of high level concepts that are both valuable and easy to use for evaluating the designs and assessment of informational artifacts, such as software applications. An example dimension is 'viscosity', which simply refers to resistance to local change. The analogy of stirring a spoon in treacle (high viscosity) versus milk (low viscosity) quickly gives the idea. Having understood the concept in a familiar context, Green then shows how the dimension can be further explored to describe the various aspects of interacting with the information structure of a software application. In a nutshell, the concept is used to examine "how much work you have to do if you change your mind" (Green, 1990, p79). Different kinds of viscosity are described, such as 'knock-on' viscosity, where performing one goal-related action makes necessary the performance of a whole train of extraneous actions. The reason for this is due to constraint density: the new structure that results from performing the first action violates some constraint, which must be rectified by the second action, which in turn leads to a different violation, and so on. An example is editing a document using a word processor without widow control. The action of inserting a sentence at the beginning of the document can have a knock-on effect whereby the user must then go through the rest of the document to check that all the headers and bodies of text still lie on the same page.

One of Green's claims about the value of cognitive dimensions is that by identifying different kinds of dimensions at a suitable level of abstraction across applications, solutions found in one domain may be applicable to similar problems found in others. Such a lingua franca of design concepts is proving to have much appeal. Various people have used and adapted the conceptual framework to determine why some interfaces are more effective than others. These include educational multimedia (e.g. Oliver, 1997; Price, 2002), collaborative writing (Wood, 1995) and various programming environments (Modugno *et al.,* 1994; Yang *et al.,* 1995). In contrast with activity theory concepts, designers and researchers, alike, who have been exposed for the first time to the dimensions have found them comprehensible, requiring not too much effort to understand and to learn how to use (Green *et al.,* 1996). Indeed, when one first encounters the 'cog dims' there is a certain quality about them that lends to articulation. They invite one to consider explicitly trade-offs in design solutions that might otherwise go unnoticed and which, importantly, can be traced to the cognitive phenomena they are derived from.

Our own approach to making the theory of external cognition applicable to design concerns (Scaife and Rogers, 1996; Rogers and Scaife, 1998) was based on an analysis of how graphical representations are used during various cognitive activities, including learning and problem-solving. Our primary objective was to explain how different kinds of graphical representations (including diagrams, animations and virtual reality) are interacted with when carrying out cognitive tasks. The *properties* and *design dimensions* that we derived from this, were intended to help researchers and designers determine which kinds and combinations of graphical representations would be effective for supporting different kinds of activities. A central property we identified is computational offloading – the extent to which different external representations vary the amount of cognitive effort required to carry out different activities. This is described further in terms of other properties, concerned with the nature of how different external representations work. We also operationalized particular design dimensions as design concepts, intended to be used at a more specific level, to guide the design of interactive representations (see figure 2). An example of a design concept is cognitive tracing, which refers to the way users are allowed to develop their own understanding and external memory of a representation of a topic by being allowed to modify and annotate it.

At the highest conceptual level, **external cognition** refers to the interaction between internal and external representations when performing cognitive tasks (e.g. learning). At the next level this relationship is characterized in terms of:

• **computational offloading** - the extent to which different external representations reduce the amount of cognitive effort required to solve informationally equivalent problems

This is operationalized in terms of the following dimensions:

• **re-representation** - how different external representations, that have the same abstract structure, make problem-solving easier or more difficult

• **graphical constraining** - this refers to the way graphical elements in a graphical representation are able to constrain the kinds of inferences that can be made about the underlying represented concept

• **temporal and spatial constraining** - the way different representations can make relevant aspects of processes and events more salient when distributed over time and space.

For each of these dimensions we can make certain predictions as to how effectively different representations and their combinations work. These dimensions are then further characterized in terms of design concepts with the purpose of framing questions, issues and trade-offs. Examples include the following:

• **explicitness and visibility** – how to make more salient certain aspects of a display such that they can be perceived and comprehended appropriately

• **cognitive tracing** – what are the best means to allow users to externally manipulate and make marks on different representations

• **ease of production** – how easy it is for the user to create different kinds of external representations, e.g. diagrams and animations

• **combinability and modifiability** – how to enable the system and the users to combine hybrid representations, e.g. enabling animations and commentary to be constructed by the user which could be appended to static representations

Figure 2: A theoretical framework of cognitive interactivity (adapted from Rogers and Scaife, 1997)

In turn, this concept provides the designer with a way of generating possible functions at the interface in a particular graphical form that supports the above. For example, Masterman and Rogers (2002) developed a number of online activities that allows children to create their own cognitive traces when learning about chronology using an interactive multimedia application. These included a drag and drop technique that allowed them to match days of the week to the deities from whom their names were derived (see figure 3).

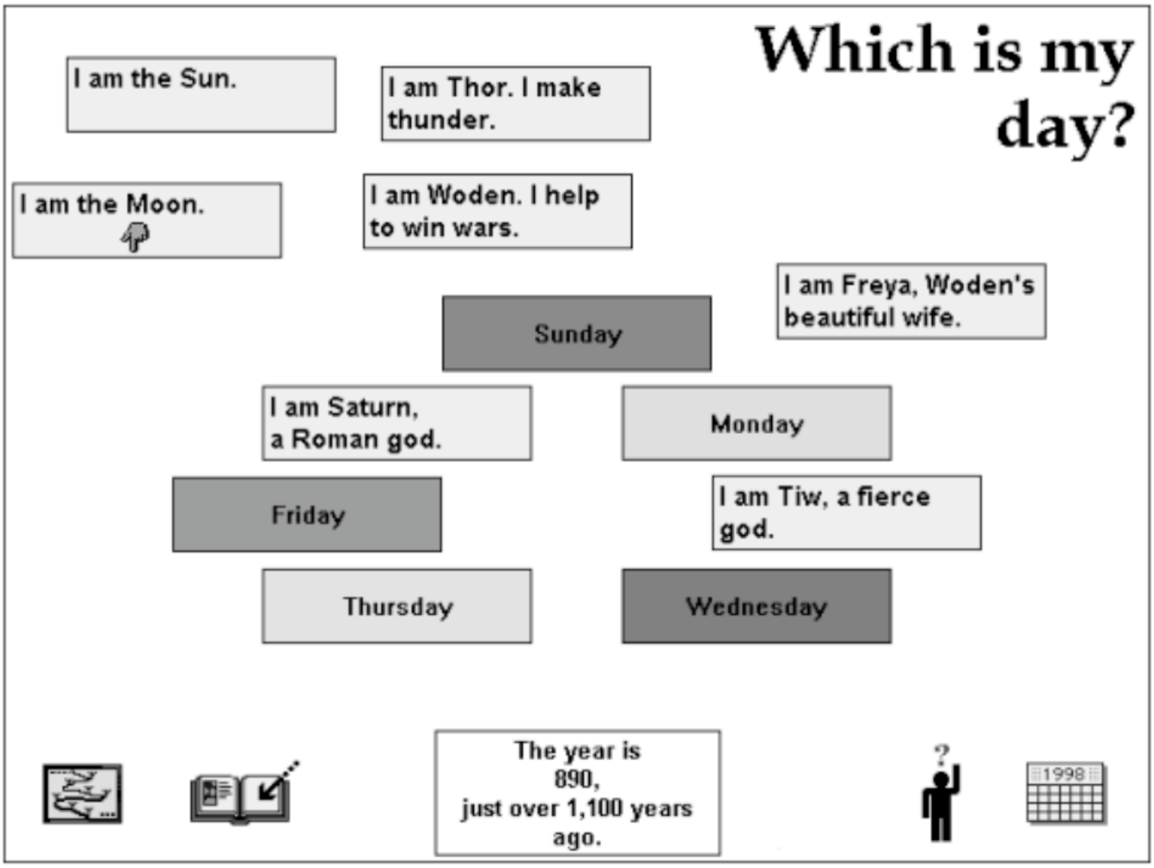

Figure 3. Example of the application of the design principle of cognitive tracing: the task is to drag each god onto the appropriate visual description of each day name. On this screen Tuesday and Sunday have already been matched to their respective deities and the mouse pointer indicates that the user can drag the statement 'I am the Moon' to a destination (i.e. Monday). (From Masterman and Rogers, 2002, p235)

So far, the set of concepts and dimensions have been most useful for deciding how to design and combine interactive external representations for representing difficult subjects, such as dynamical systems in biology, chronology in history, the working of the cardiac system and crystallography (e.g. Gabrielli *et al.,* 2000; Masterman and Rogers, 2002; Otero, 2003; Price, 2002). Sutcliffe (2000) has also shown how he used the theory to inform the design of multimedia explanations. More recently, we have used the approach in work settings, to inform the design of online graphical representations that can facilitate and support complex distributed problem-solving (Scaife *et al.,* 2002; Rodden *et al.,* 2003).

One of the main benefits of our approach is the extent to which the core properties and design dimensions can help the researcher select, articulate and validate particular forms of external representation in terms of how they can support the activity being designed for. Its emphasis on determining the optimal way of structuring and presenting interactive content with respect to the cognitive effort involved, is something we would argue other theoretical approaches, like activity theory and the ecological approach, do not do, since their focus has been more on elucidating the nature of existing problems. In sum, the way theory has been used to inform the cognitive and design dimensions approaches, is largely *generative*.

The distributed cognition approach

The distributed cognition approach was developed by Hutchins and his colleagues in the mid to late 80s and proposed as a radically new paradigm for rethinking all domains of cognition (Hutchins, 1995). It was argued that what was problematic with the classical cognitive science approach was not its conceptual framework *per se,* but its exclusive focus on modeling the cognitive processes that occurred within one individual. Alternatively, Hutchins argued, what was needed was for the same conceptual framework to be applied to a range of cognitive systems, including socio-technical systems at large, (i.e. groups of individual agents interacting with each other in a particular environment). Part of the rationale for this extension was that, firstly, it was assumed to be easier and more accurate to determine the processes and properties of an 'external' system – since they can arguably, to a large extent, be observed directly in ways not possible inside a person's head – and, secondly, they may actually be different and thus unable to be reduced to the cognitive properties of an individual. To reveal the properties and processes of a cognitive system requires doing an ethnographic field study of the setting and paying close attention to the activities of people and their interactions with material media (Hutchins, 1995). Similar to the external cognition approach, these are conceptualized in terms of "internal and external representational structures" (Hutchins, 1995, p135). It also involves examining how information is propagated through different media in a cognitive system.

The distributed cognition approach has been used primarily by researchers to analyze a variety of cognitive systems, including airline cockpits (Hutchins and Klausen, 1996; Hutchins and Palen 1997), air traffic control (Halverson, 1995), call centers (Ackerman and Halverson, 1998), software teams (Flor and Hutchins, 1992), control systems (Garbis and Waern, 1999) and engineering practice (Rogers, 1993, 1994). One of the main outcomes of the distributed cognition approach is an explication of the complex interdependencies between people and artifacts in their work activities. An important part of the analysis is identifying the problems, breakdowns and the distributed problem-solving processes that emerge to deal with them. In so doing, it provides multi-level accounts, weaving together "the data, the actions, the interpretations (from the analyst), and the ethnographic grounding as they are needed" (Hutchins and Klausen, 1996, p.19). For example, Hutchins' account of ship navigation provides several interdependent levels of explanation, including how navigation is performed by a team on the bridge of a ship; what and how navigational tools are used, how information about the position of the ship is propagated and transformed through the different media and the tools that are used.

As a theoretical approach, it has received considerable attention from researchers in the cognitive and social sciences, most being very favourable. However, there have been criticisms of the approach, mainly as a continuation of an ongoing objection to cognitive science as a valid field of study and, in particular, the very notion of cognition (e.g. Button, 1997). In terms of its application in HCI, Nardi (1996, 2002) has been one of the most vociferous in voicing her concerns about its utility in HCI. Her main criticism stems from the need to do extensive field work before being able to come to any conclusions or design decisions for a given work setting. Furthermore, she points out, that compared with Activity Theory (which she is a strong advocate of), there is not a set of interlinked concepts that can be readily used to pull things out from the data. In this sense, Nardi has a point: the distributed cognition approach is much harder to apply, since there is not a set of explicit features to be looking for, nor is there a check-list or recipe that can be easily followed when doing the analysis. It requires a high level of skill to move between different levels of analysis; to be able to dovetail between the detail and the abstract. As such it can never be viewed as a 'quick and dirty' prescriptive method. The emphasis on doing (and interpreting) ethnographic fieldwork to understand a domain, means that at the very least, considerable time, effort and skill is required to carry out an analysis.

Where the distributed cognition framework can be usefully applied to design concerns, is in providing a detailed level of analysis which can provide several pointers as to how to change a design (especially forms of representation) to improve user performance, or, more generally, a work practice. For example, Halverson (2002) discusses how in carrying out a detailed level of analysis of the representational states and processes involved at a call center, she was, firstly, able to identify why there were problems of coordination and, secondly, determine how the media used could be altered to change the representational states to be more optimal. Hence, design solutions can start to emerge from a detailed level of analysis because the nature of the descriptions of the cognitive system are at the same level as the proposed design changes. Moreover, as Halverson (2002) points out, this contrasts with using an Activity Theory framework, because the outcome of doing an analysis using AT concepts, is at a higher level that does not map readily onto the level required for contemplating design solutions. Hence, her argument is that it is because of rather than in spite of, the low level nature of the analysis that can be most useful at revealing the necessary information to know how to change a design, when it has been identified as being problematic.

More generally, the distributed cognition approach can inform design by examining how the form and variety of media in which information is currently represented might be transformed and what might be the consequences of this for a work practice. Partially in response to the criticism leveled at the difficulty of applying the distributed cognition approach, Hutchins and his colleagues (Hollan *et al.,* 2000) have set an agenda for how it can be used more widely within the context of HCI. They propose it is well suited both to understanding the complex networked world of information and computer-mediated interactions and for informing the design of digital work materials and collaborative workplaces. They suggest a comprehensive methodological framework for achieving this – albeit at this stage a somewhat ambitious and complex programme. The way theory has been applied from the DC approach, has been largely *descriptive* and to a lesser extent *generative*; providing a detailed articulation of a cognitive system, and in so doing, providing the basis from which to generate design solutions.

The situated action approach

The situated action approach has its origins in cultural anthropology (Suchman, 1987). Its rationale is based on the proposed need for "accounts of relations among people, and between people and the historically and culturally constituted worlds that they inhabit" (p71, ibid). A main goal is to "explicate the relationship between structures of action and the resources and constraints afforded by physical and social circumstances" (p179, ibid). This is accomplished by studying "how people use their circumstances to achieve intelligent action (...) rather than attempting to abstract action away from its circumstances"(p. 50, ibid). Furthermore, it views human knowledge and interaction as being inextricably bounded with the world: "one cannot look at just the situation, or just the environment, or just the person", since to do so, "is to destroy the very phenomena of interest" (Norman, 1993, p. 4). Hence, its epistemological stance is the very antithesis of the approaches we have described so far: resisting any form of theoretical abstraction.

The method used is predominantly ethnographic (i.e. carrying out extensive observations, interviews and note-taking of a particular setting). Typically, the findings are contrasted with the prescribed way of doing things, i.e. how people ought to be using technology given the way it has been designed. For example, one of the earliest studies, using this approach was Suchman's (1983) critique of office procedures in relation to the design of office technology. Her analysis showed how there is a big mismatch between how work is organized in the process of accomplishing it in a particular office and the idealized models of how people should follow procedures that underlie the design of office technology. Simply, people do not act or interact with technology in the way prescribed by these kinds of models. Instead, Suchman argues that designers would be much better positioned to design systems that could match the way people behave and use technology if they began by considering the actual details of a work practice. The benefits of doing so could then lead to the design of systems that are much more suited to the kinds of interpretative and problem-solving work that are central to office work.

In her later, much cited, study of how pairs of users interacted with an expert help system – intended as a help facility for using with a photocopier – Suchman (1987) again stresses the point that the design of such systems would greatly benefit from analyses that focus on the unique details of the user's particular situation – rather than any preconceived models of how people ought (and will) follow instructions and procedures. Her detailed analysis of how the expert help system was unable to help users in many situations where they got stuck, highlights once more the inadequacy of basing the design of an interactive system primarily on an abstract user model. In particular, her findings showed how novice users were not able to follow the procedures, as anticipated by the user model, but instead engaged in on-going, situated interaction with the machine with respect to what they considered at that moment as an appropriate next action.

These kinds of detailed accounts provide much insight into how technology is actually used by people in different contexts, which is often quite different from the way the technology was intended to be used. Moreover, their influence on the field has become quite pervasive. Several researchers have reported how the situated action approach has profoundly changed the way they think about how they conceptualise and develop system architectures and interface design (e.g. Button and Dourish, 1996; Clancey, 1997). More generally, Suchman has been one of the most frequently cited authors in the HCI literature. The approach has also become part of designer's talk; concepts of 'situatedness' and 'context' often being mentioned as important to design for. Hence, the situated action approach has, arguably, had a considerable influence on designers. Nowadays, it is increasingly common for designers and others to spend time 'in the field' understanding the context and situation they are designing for before proposing design solutions (Bly, 1997). For example, large corporations like Microsoft, Intel and HP, have recently begun to make claims about the benefits of this approach in their online promotional blurb, e.g.,

> "Field studies open our eyes to how regular people, unguided, use their PC and the Web, as well as specific products and features we design. We use the resulting information to guide us in the redesign and enhancement of our products to reflect how people want to use them." (Microsoft, 2002, p4)

One of the main criticisms of the situated action approach, however, is its focus on the 'particulars' of a given setting, making it difficult to step back and generalize. Similar to the criticism leveled at about doing field studies using the distributed cognition approach, Nardi (1996) exclaims how in reading about the minutiae of a particular field study “one finds oneself in a claustrophobic thicket of descriptive detail, lacking concepts with which to compare and generalize” (p.92). It seems those who are used to seeing the world through abstractions find it hard to conceptualise and think about design at other levels of detail.

Others have taken on board this criticism and have attempted to draw some core abstractions from the corpus of field studies, that have been concerned with situatedness and context. Most notable, is Hughes *et al’s* (1997) framework developed specifically to help structure the presentation of ethnographic findings in a way that was intended to act as a bridge between fieldwork and ‘emerging design decisions’. The abstractions are discussed in terms of three core dimensions (a similar method of abstraction to the external cognition approach). As such, they are intended to orient the designer to thinking about particular design problems and concerns in a focused way, that in turn can help them articulate why a solution might be particularly helpful or supportive.

Contextual design (Beyer and Holzblatt, 1998) is another approach that was developed to deal with the collection and interpretation of ethnographic findings and to use these to inform the design of software. In contrast to the dimensions approach described above, it is heavily prescriptive and follows a step-by-step process of transforming data into a set of abstractions and models. Part of its attraction is because of its emphasis on heavyweight conceptual scaffolding, providing the user with a recipe to follow, and various ‘forms’ to fill in and use to transform findings into more formal structures. However, in so doing, its relationship with the situated action approach is inevitably divorced, since its focus is more on how to progress layers of abstractions rather than bridging analysis and design through examining the detail of each.

In sum, the influence of the situated action approach on HCI practice has been divergent. On the one hand, it’s contribution has been *descriptive*, providing accounts of working practices, and on the other, it has provided a *backdrop* from which to talk about high level concepts, like context. It has also inspired and led to the development of analytic frameworks and core dimensions.

The ethnomethodological approach

Ethnomethodology is an analytic framework, that was originally developed as a reaction against the traditional approaches in sociology, which were largely top-down theories geared towards identifying invariant structures (Garfinkel, 1967; Garfinkel and Sacks, 1970). Such external points of view of the world were considered not at all representative of the actual state of affairs. In this sense, it adopts an anti-theoretical stance and is very outspoken about its epistemological origins. Alternatively, the ethnomethodologists argue for a bottom-up approach, whereby working practices are described from the practical accomplishment of the people (Anderson, 1994). To achieve this, the approach adheres to a rigorous descriptive programme, that accounts for members’ (sic) working practices.

Similar to the situated action and distributed cognition approaches, it has been used to explicate the details of various work practices through which actions and interactions are achieved. It has been popularized mainly by British sociologists, who have used it to analyze a number of workplace settings, the most well known being a control center in the London Underground (Heath and Luff, 1991) and air traffic control (Bentley *et al.,* 1992). These accounts of work practices are presented largely as thick (Geertz, 1993) descriptions. By this it is meant extensive and very detailed accounts. In the same vein as the situated action based ethnographies, the detailed accounts have proved to be very revealing, often exposing taken for granted working practices, which turn out to be central to the efficacy of how a technological system is being used.

To show how these accounts might be useful for the design of technology and work, 'design implications' are typically teased out of them, but in, unfortunately, a somewhat superficial manner. The problem of requiring ethnomethodologists to venture into this unfamiliar territory – namely, offering advice for others to follow – is that it typically ends up being little more than a cursory set of tepid guidelines. Part of the reason for this uncomfortable state of affairs is that the ethnomethodologists simply feel ill-equipped to offer advice to others, whose very profession is to design – which clearly theirs is not. Their role is regarded as descriptive not prescriptive (Cooper, 1991). For example, in one study Anderson *et al.* (1993) provided a very detailed and insightful descriptive account of an organization's working practice. Following this, they outlined four brief 'bullet-point' guidelines. One of these is that designers need support tools that take up a minimal amount of their time and that such tools should be adaptive to the exigencies of changing priorities. Such an observation is stating the obvious and could have easily been recognized without the need of a detailed field study. It is not surprising that this form of abstracting from detailed field studies was derided; "most designers know the former only too well and desire the latter only too much" (Rogers, 1997, p68).

Recognizing the dilemma confronting ethnomethodologists entering the field of HCI, resulted in a rethinking of what else they could offer in addition to the thick descriptions (Geertz, 1993) and token nuggets, that could be perceived to be more useful to design concerns. Ironically, it was the core set of social mechanisms, that were written about by the founders of ethnomethodology, that provided them with a way forwards. Button and Dourish (1996), for example, discuss how the high level socially-based concepts of practical action, order, accountability and coordination could be potentially of more value to designers. Furthermore, they proposed that ethnomethodologists and designers could greatly benefit by trying to see the world through each other's perspective: "design should adopt the analytic mentality of ethnomethodology, and ethnomethodology should don the practical mantle of design" (p. 22). It was suggested that this form of synergism could be achieved through system design taking on board 'generally operative processes' like situatedness, practical action, order and accountability, whilst ethnomethodology could take on system design concepts like generalization, configuration, data and process and mutability. To show how this forging of theory might work, a hypothetical example of two different questions was given that might be asked when designing a new system. Rather than ask "what are the implications of this ethnomethodological account of the work of hotel receptionists for the design of a booking system" (p.22) they suggest a more insightful question might be "what are the implications of the operation and use of member categories for questions of individuality and grouping in software systems?" (p22). However, whilst highlighting a more specific requirement for a system, it is difficult to imagine designers (or others) ever becoming sufficiently versed in this kind of discourse (referred to as 'technomethodology') to talk about design issues to each other in this way. Moreover it is privileging a form of academic 'hybrid' talk, that to most 'plain' folk can seem arcane and cumbersome. Some might argue, however, that as if with any new set of concepts, once time and effort has been spent learning how to use them, then their benefits will accrue. Having learnt the new way of talking, then designers and others would be able to extend their discourse and articulate design problems in a more illuminating and explicit way. This indeed may prove to be the case and it is the argument put forward by Green (1989) in his exposition of the vocabulary of cognitive dimensions.

However, one cannot help thinking that the ethnomethodologically-based concepts will prove to be much harder to learn and use in the context of a design space than the likes of viscosity, cognitive offloading and affordances, which designers have found useful and relatively easy to use. It is the like the difference between learning to speak French and Norwegian as a second language.

In sum, the ethnomethodological approach, like the situated approach, began with providing *detailed descriptions* of work practices – assuming this was a significant contribution for HCI and has more recently, sought alternative ways of informing design, through providing *a linga franca*, comprising a set of core concepts.

Hybrid and overarching theoretical approaches

Besides importing and developing individual approaches in HCI, several researchers have tried to synthesize concepts from different theories and disciplines. A main rationale for this strategy is to provide more extensive frameworks than if they were to import concepts arising from only one discipline. In attempting to articulate relevant concerns, Star (1996) for example, has drawn parallels between different strands of different theories. In one instance, she has looked at similarities between activity theory and symbolic interactionalism (originating from American pragmatism) with a view towards forging better links between them. More ambitiously, Pirolli and Card (1997) have reconceptualized a particular form of human-computer interaction, namely searching for and making sense of information, using a variety of concepts borrowed from evolution, biology and anthropology together with classical information processing theory: "we propose an information foraging food-theory (IFT) that is in many ways analogous to evolutionary ecological explanations of food-foraging strategies in anthropology and behavior ecology" (p. 5). They describe searching strategies in terms of making correct decision points, which are influenced by the presence or absence of 'scent'. If the scent is strong enough, the person will make the correct choices; if not they will follow a more random walk. Their approach is replete with such metaphors, re-describing activities in terms of more concrete everyday experiences. In so doing, it has enabled the authors to rethink the field of information visualization, informing the development of new kinds of graphical representations and browsing tools.

Perhaps the most ambitious attempts at developing theory for HCI are the overarching frameworks that attempt to integrate multiple theories at different levels of analysis. For example, Mantovani's (1996) eclectic model for HCI integrates a wide range of concepts and research findings that have emerged over the last 10 years, from computer supported-cooperative work (CSCW), computer mediated communication (CMC) and distributed artificial intelligence (DAI). The outcome is a three level conceptual model of social context, that combines top-down with bottom-up approaches to analyzing social norms and activities. Likewise, Barnard *et al's* (2000) 'Systems of Interactors' theoretical framework, draws upon several overlapping layers of macrotheory and microtheory. Which level of theory is relevant depends on the nature of the problem being investigated.

A problem with integrating quite different theories and ontologies, however, is that it makes it very difficult to know what frames of reference and axioms to use for a given problem space. Furthermore, it can be quite unwieldy to juggle with multiple concepts, constraints and levels when analyzing a problem space and/or designing a system. It seems only the researchers, themselves, who have developed the 'grand' theories, are able to use them.

In sum, a main objective of developing hybrid and overarching frameworks for HCI is to provide a more extensive, interdisciplinary set of concepts, from which to think about the design and use of interactive systems. A commonly reported benefit of pursuing this is that it allows one to break away from the confines of a single discipline, and in so doing, evolve new ideas, concepts and solutions. In this sense the theory can serve a *formative* and *generative* role for design. Certainly, one of the benefits of juxtaposing and interweaving different concepts from different traditions is that it can create new perspectives and ways of thinking about a problem space. The danger of this approach, however, is that resultant frameworks can simply be too unwieldy to apply to specific design concerns, especially if the designers/researchers are not au fait with the ideas originating from the parent disciplines. As such, they are likely to suffer from the toothbrush syndrome:

> "Ernest Hilgard used to grumble about psychology that if you develop a theory it's like your toothbrush, fine for you to use but no one else is very interested in using it." Grudin (2002, ChiPlace online forum)

The practitioner's perspective

My critique and overview of the role of theories have recently been imported and developed in HCI has so far been based primarily on a review of the HCI literature. Here, I consider the practitioner's perspective of the role of theory in practice, based on what they report they use in their work. By practitioner, I mean people who work in industry and are in the business of researching, designing and evaluating products (e.g. interaction designers, information architects, usability experts). The intention of this section is to highlight what they think the role of theory is in HCI and their perceived needs for it in the kind of work they do. It presents some provisional findings from a small survey carried out by myself and summarizes the findings of another survey that was carried out in Sweden, by way of comparison (Clemmensen and Leisner, 2002).

The initial survey I carried out was designed as an online questionnaire and was sent to 60 practitioners, from the UK and the US. Rather than carry out in-depth interviews with a relatively small number of people (the more widely accepted method for doing survey work) I wanted to get a larger set of 'quick and dirty' responses from a range of people working in quite different organizations. To achieve this, I adopted the pyramid approach; sending out the questionnaire to a range of people I knew working in large corporations (e.g., IBM, Microsoft, HP, Logica, Motorola), medium-sized design companies (e.g., VictoriaReal), and small interaction design consultancies (e.g., Swim) and asking them to fill it in and also forward it on to their colleagues. A total of 34 people responded, of which 12 classified themselves as doing mainly design, 10 classified themselves as doing mainly research, 4 doing a mix of activities, 4 doing mainly production work and 4 doing mainly usability evaluation. Although the number of respondents is still relatively small, the spread is sufficiently broad, to get a sample of views.

The questionnaire asked a number of questions about their current practice and in particular whether they had heard about the theories presented in the previous section, and if they had used any of the concepts and analytic frameworks in their work. The respondents were first asked what methods they used in their work. Nearly all replied that they used a range of design methods, including scenarios, storyboards, sketching, lo-tech and software prototyping, focus groups, interviews, field studies and questionnaires and use cases. None of them used predictive modeling methods, like GOMS, while a few used software engineering methods (8%), experiments (10%), contextual design (10%) or guidelines (5%).

The combination of methods used by the respondents indicates that there is much gathering of information and requirements in their work. This suggests that there is a need for it to be interpreted and analyzed in some way. When asked what they use to interpret their findings, however, 85% of the respondents said that they relied mainly on their own intuition and experience. The few who did say they used theory, said they did so only occasionally. The theories used were either their own adaptation, distributed cognition, or grounded theory. Interestingly, this lack of use of recently imported theoretical approaches contrasted markedly with the knowledge that the respondents said they have about them. Indeed, many of the respondents claimed to be familiar with most of the approaches mentioned in the previous section (see figure 5). Thus, it seems that while many practitioners may be familiar with the approaches that have been promoted in HCI, very few actually use them in their work, and even then only sporadically.

Part of the problem seems to be the gap between the demands of doing design and the way theory is conceptualised, as commented on by respondent 14 (who described himself as a designer): "most current HCI theory is difficult for designers to use and generally too theoretical to be relevant to a practical human focused solution developed in the timeframe of a design project."

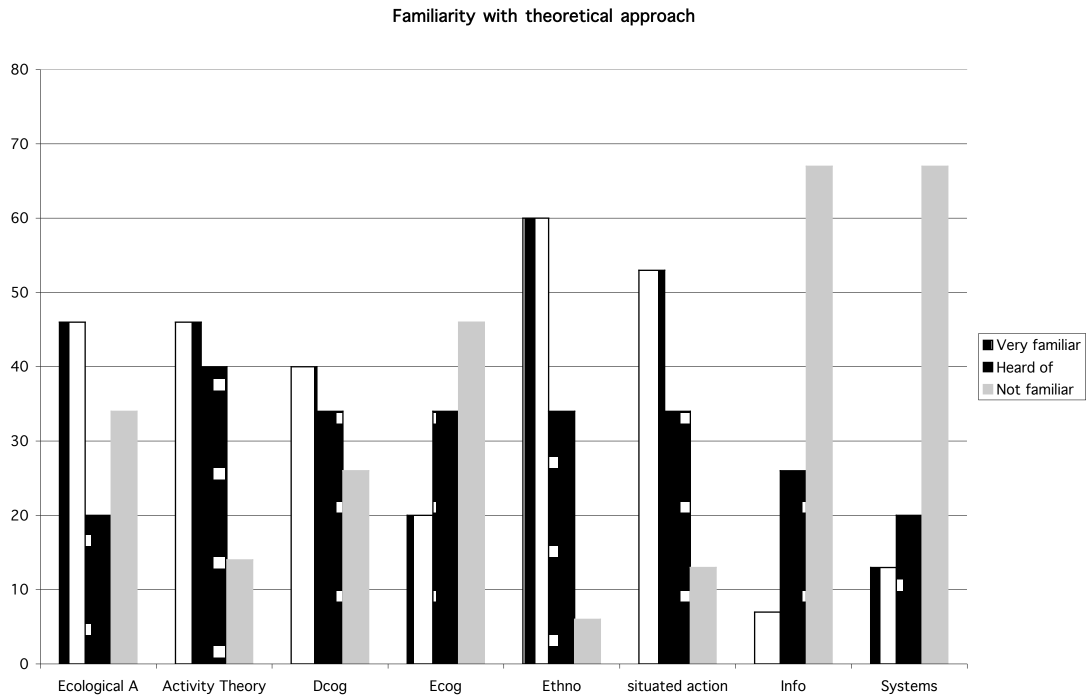

Figure 5. Respondent's familiarity (as a percentage of total responses) with theoretical approaches (left column- very familiar, middle column- heard of, and last column- not familiar with.

In contrast to the lack of uptake of recent theoretical approaches as analytic frameworks, the concepts derived from them were found to be more commonly used by the respondents when talking with others about their work. Many said they used the concepts of affordances (75%), context (80%), awareness (65%), situatedness (55%) and cognitive offloading (45%). Concepts that were less used were ecological constraints (25%), cognitive dimensions (15%) and propagation of representational states (10%).

Thus, it seems that a number of concepts, especially those derived from the situated action approach are commonly used as part of the discourse with work colleagues.

When asked whether they found it difficult to express ideas about a project to others in their group (or clients), the opinions were divided between those replying, “all the time” (30%), those responding “some of the time” (45%) and those saying “no problem” (25%). The findings suggest, therefore, that over 70% of respondents have trouble communicating ideas with others. When asked whether they would like a better set of terms and concepts to use, 50% of the respondents said yes, 35% said not sure and 15% said they were happy with the way they communicated. Interestingly, when asked whether there was a need for new kinds of analytic frameworks, an overwhelming 92% said yes. When asked what else they would find useful, many replied that there was a need for existing frameworks to be better explained. For example:

Respondent 6 (designer) asked for a “framework for effectively communicating with clients...a common language between designer and client seems to be lacking.”

Respondent 10 (designer) asked for “more support for guidance in applying the existing frameworks.”

Respondent 22 (consultant) asked for “better ways of talking about existing frameworks…better ways of talking about how situated action or ethnomethodology (or any other theory) informs the practice I use in a way that makes sense to a person unfamiliar with the underlying theory.”

This small survey has revealed that even though practitioners are familiar with many of the recent theoretical approaches that have been imported into the field of HCI, they don’t use them in their work, because they are too difficult to use. Moreover, it is not that they don’t find them potentially useful, but that they do not know how to use them. This contrasts with Bellotti’s (1988) study, where she suggested that one of the main reasons why designers did not use any of the HCI techniques at the time was because they had no perceived need for them, regarding them as too time-consuming to be worthwhile. A frequently cited complaint was that they wanted more guidance and ways of communicating about them to others.

In a more extensive survey of Danish usability professionals, researchers and designers (120 in total), Clemmensen and Leisner (2002) asked their respondents to consider the relationship between the publicity different theories received in the HCI community and how applicable they were. The range of theories that the respondents were asked to judge was similar to those discussed in this chapter. Similar to my study, they found that most of the respondents were interested in different theories, favoring one or two kinds. In contrast to my findings, however, they found that over 50% of the Danish usability professionals said that they used at least one theory in their investigations. One of the reasons for this contrast in results may have to do with the sampling: the Danish usability specialists were more similar; all young, having less than 5 years experience and all having a PhD from the social sciences, with over half having written about HCI issues. In contrast, my sample of respondents covered a much wider age span, and more diverse cultural, educational and professional backgrounds. The Danish professionals were all part of an online community and hence could be regarded as a self-selecting group. The questions asked were also worded differently, inviting the respondents to back their claims as to why they found theory useful (e.g. one respondent said, “I want my work to have a theoretical basis, to have the framework for understanding and assurance of a methodology that helps me explain the results of investigation". In my case, I asked them if they found theory useful and in what ways.

In sum, the findings from the two surveys of practitioner's use of theory indicate that they are interested in using theory in their work. What they can use they do use: for example, they use several of the *concepts* derived from the theories in their discourse. However, from my study it seems that often practitioners do not know how to apply the much harder to use *analytic frameworks* to the specifics of the projects they are involved in (e.g. the field data they gather). Part of the dilemma facing practitioners is the pressure they are under to solve problems quickly and under 'deadline' while at the same time wanting to ground their work, theoretically. As argued earlier, to do justice to many of the analytic frameworks that have been developed in HCI, based on theory, one needs, firstly, a good apprenticeship in it, and secondly, the time, patience and skill to competently carry out a detailed analysis. Given that many practitioners are unlikely to satisfy both requirements, it seems that the analytic frameworks, such as Activity Theory, and distributed cognition will continue to remain out of reach. Alternatively, approaches to bridging the gap between theory and practice, that are more lightweight and accessible may prove to have more utility.

Discussion

My overview of the earlier and more recent theoretical approaches imported, developed and applied in HCI has shown that there is a difference between how they have been used in the field. Primarily, the way in which theory was used by the earlier approaches was :

- **informative** (providing useful research findings)
- **predictive** (providing tools to model user behavior)
- **prescriptive** (providing advice as to how to design or evaluate)

The way theory has been used in the newer approaches is more diverse:

- **provide descriptive accounts** (rich descriptions)
- **be explanatory** (accounting for user behavior)
- **provide analytic frameworks** (high level conceptual tool for identifying problems and modeling certain kinds of user-interactions)
- **be formative (**provide a lingua franca; a set of easy to use concepts for discussing design)
- **be generative** (provide design dimensions and constructs to inform the design and selection of interactive representations).

Hence, there appears to have been a move away from providing predictive and prescriptive approaches towards developing more analytic and generative approaches. One of the most significant contributions has been to provide more extensive and often illuminating accounts of the phenomena in the field. A further contribution has been to show the importance of considering other aspects besides the internal cognitive processing of a single user – notably, the social context, the external environment, the artifacts and the interaction and coordination between these during human-computer interactions. All of which can help towards understanding central aspects of the diffuse and boundless field that HCI has become.

We now have a diverse collection of accounts and case studies of the intricate goings on in workplace and other settings (e.g. Plowman et al, 1995). An eye for detail, resulting in an analysis of the normally taken-for-granted actions and interactions of people in particular contexts, has shown us the instrumental role of a range of social and cognitive mechanisms. Analogous to the literary works of Nicholson Baker and Ian McEwan – that both offer lucid and intimate accounts of the mundane that enable us to perceive everyday occurrences and artifacts in a new light – many of the detailed ethnographically-informed accounts of situated human-computer interactions have opened our eyes to seeing the world of technology use quite differently. In turn, this can lead us to thinking about the design and redesign of technologies from quite different perspectives.

Another significant development is the pervasive use of a handful of high level concepts derived from the new approaches. These have provided different ways of thinking and talking about interaction design. As the two surveys revealed, practitioners are aware of various concepts, like situatedness, context and awareness, which they use when talking with others during their work. Clearly, such concepts provide a way of articulating current concerns and challenges, that go beyond the single user interface.

In an attempt to be more applied, many of the new approaches have sought to construct conceptual frameworks rather than developing fully-fledged theories in the scientific Popperian tradition. Frameworks differ from theories in that they provide a set of constructs for understanding a domain rather than producing testable hypotheses (Anderson, 1983). The value of adopting this more relaxed research strategy is that it enables a broadening of scope – something which has now become widely recognized as having been a necessary step for developing better accounts of human-computer interaction. However, ironically, it appears that the analytic frameworks developed for use in HCI are not that accessible or easy to use. Designers, consultants, producers and others involved in the practice of interaction design are much less likely to have the time to develop and practice the skills necessary to use the analytic frameworks (e.g. carry out an activity theory or distributed cognition analysis) – echoing a similar complaint that was often made about using cognitive task analytic tools (Bellotti, 1988). This raises the question as to whether such analytic frameworks are an appropriate mechanism for practitioners to use in their work, or, whether the community should accept that they are simply too hard, requiring too much time and effort to use, and should be left for those doing research. If the latter is the case, then can we find other ways of translating theory-based knowledge that is easier to use and fits in with the perceived needs of practitioners?

In the next section I discuss the reasons why theoretically-informed tools appear to be finding it difficult to infiltrate actual design practice, and then in the final part I propose how this gap can be more effectively bridged.

Why are alternative theories problematic in practice?

When the 'second' generation of alternative approaches began to be introduced into the field of HCI there was considerable skepticism as to what they had to offer of practical value that would persuade designers to take them on board. For example, in a review of Bødker's (1989) book on Activity Theory and HCI, Draper (1992) notes how her application of concepts from Activity Theory to HCI do not add to the existing set of ideas about design, nor convince newcomers about the potential of Activity Theory. Nardi (1996) has also been critical of the value of and methodological positions adopted by the distributed cognition and situated action approaches. So why have these attempts not been well received within parts of the HCI community?

There are several reasons why the new approaches have yet to make a more marked impact on the *process* of interaction design (as opposed to just becoming part of the body of HCI knowledge). Firstly, it must be stressed that it is foolish to assume or hope that theories "do design" however much the proponents of the theoretical approach would like (Barnard and May, 1999). Their input to the design process can only ever really be indirect, in the form of providing methods, concepts, frameworks, analytic tools and accounts. A theory cannot provide prescriptive guidance in the sense of telling a designer what and how to do design. The contribution of any theory must be viewed sensibly and in the context of its role in the design process at large. Designers already have an armory of practical methods and techniques available to them to use (e.g. prototyping, heuristic evaluation, scenario-based design). For this reason, the value of theory-informed approaches must be seen in relation to current design practice.

Secondly, more time is needed to allow a complete theory/design cycle to mature (e.g. Plowman *et al.,* 1995). It may take several more years before we see more success stories being reported in the literature – just as it took several years after GOMS was developed before its value in a real work setting was able to be reported. Such case studies could be set up as exemplars of good practice for designers to learn lessons from in how to apply the approach. The use of case studies as a way of explaining an approach is much more common in design.

Thirdly, as emphasized throughout this chapter, considerable time, effort and skill are required by many of the approaches to understand and know how to use them. In particular, many require ethnographic field work to be carried out as part of the approach. Knowing how to 'do' ethnography and to interpret the findings in relation to a theoretical framework (e.g. activity theory, distributed cognition) is a highly skilled activity, that requires much painstaking analysis. It is hard to learn and become competent at: many a student in HCI, has been attracted by the ethnographic approach and the theoretical framework of distributed cognition or activity theory, only to find themselves, in the midst of a field study, surrounded by masses of 'raw' video data without any real sense of what to look for or how to analyze the data in terms of, say, 'propagation of representational state across media' or 'actions, operations and activities'. Moreover, analytic frameworks, like Activity Theory, are appealing because of their high level of rhetorical force and conceptual scaffolding; whereby the act of naming gives credence to the analysis.

More generally, is the problem that there is little consensus as to what contribution the various approaches can or should make to interaction design. The transfer vehicles that became the standard and generally accepted 'deliverables' and 'products' for informing design during the 80s (e.g. design principles and guidelines, style books, predictable and quantifiable models) tend now to be regarded as lass appropriate for translating the kinds of analyses and detailed descriptions that recent theoretical approaches that have been imported into HCI have to offer. There is also more reticence towards the rhetoric of compassion (Cooper, 1991) and forcing one's own views of what needs to be done on another community. So what is replacing this form of design guidance?

The analytic frameworks that are being proposed, like Activity Theory, suffer from being under-specified, making it difficult to know whether the way one is using them is appropriate and has validity. This contrasts with the application of earlier cognitive theories to HCI, where the prescribed route outlined by the scientific method was typically followed (i.e. make hypotheses, carry out experiment to test them, determine if hypotheses are supported or repudiated, develop theory further, repeat procedure). Without the rigor and systematicity of the scientific method at hand, it is more difficult for researchers and designers, alike, to know how to use them to best effect or whether what they come up with can be validated.

A further problem from the designer's and researcher's perspective, is that there is now a large and ever increasing number of theoretical approaches vying with each other, making it more difficult for them to determine which is potentially most useful for them or, indeed, how to use one with respect to their own specific research or design concerns. Such a confusing state of affairs has been recognized in the HCI community and one or two attempts have been made to synthesize and make sense of the current melee of approaches. For example, Nardi (1996) sought to compare and contrast selected approaches in terms of their merits and differences for system design. However, given that the various approaches have widely differing epistemologies, ontologies and methods – that are often incommensurable – such comparative analyses can only ever really scratch the surface. There is also the problem that this kind of exercise can end up like comparing apples and oranges – whereby it becomes impossible, if not illogical to judge disparate approaches (cf. Patel and Groen, 1993). Championing one theoretical approach over another, without recourse to Popper's scientific paradigm to back up one's claims, often ends up being a matter of personal preference, stemming from one's own background and values as to what constitutes good design practice or research. That is not to say that one cannot highlight the strengths and problems of a particular approach and show how others have used it. Indeed, that is what I have attempted to do here and which Fitzpatrick (2003) in her overview of the CSCW literature has sought to do.

Another central issue that was highlighted in the chapter was the difference between approaches that provide more detailed accounts of human-computer interactions within the historical/socio-cultural and environmental contexts in which they occur and approaches that draw out abstractions, generalisations and approximations. The unit and level of analysis which is considered appropriate depends on the purpose of the analysis. 'High level' abstractions have been the sin que non of scientific theories, particular those concerned with making hypotheses and predictions. 'Low level' descriptions are the bread and butter of more sociologically-oriented accounts of behavior. Both can be informative for HCI and feed into different aspects of the design process. However, it requires better clarification as to how the two can be dovetailed and used together rather than being viewed as always being incommensurate.

The way forward: new mechanisms for using theory

At a general level, we need to consider the direction and role that theory should be moving towards in the field of HCI and the practice of interaction design. Part of this requires being clearer about the way theories can (or can not) be used for. In particular, there needs to be a better exposition of how theory can be used in *both* research and design. Can they serve multiple and expanding purposes e.g. as (i) explanatory tools, (ii) predictive tools, (iii) providing new concepts for the purpose of developing a more extensive design language and (iv) providing tools for use in the design process or would it be clearer and more useful for an approach to focus on only one of these contributions? Shneiderman (2002b) has suggested that there at least five kinds of theories we should be aiming for and using in HCI. These are:

- descriptive - in the sense of providing concepts, clarifying terminology and guide further inquiry

- explanatory – in the sense of explicating relationships and processes
- predictive – enabling predictions to be made about user performance
- prescriptive – providing guidance for design
- generative – in the sense of enabling practitioners to create or invent or discover something new.

The roles suggested here overlap with the types we identified earlier. These seems to be a consensus, therefore, that theory can and should be used more eclectically in HCI. One of the problems of trying to use theory for multiple purposes, however, is that it can be difficult to satisfy the demands that each requires. In particular, it can be problematic to adhere to both theoretical adequacy (i.e. that accounts are representative of the state of affairs) and also demonstrate transferability (i.e. that ideas, concepts and methods derived from the theoretical framework can be communicated and taken-up, resulting in the design and implementation of better technologies). Remaining faithful to the epistemological stance of a theoretical approach can make it difficult, if not impossible, to then provide a framework for applied concerns. Conversely, the approach can no longer adhere to the epistemology of the original theory when taking design concerns into account. A problem of doing this, as we saw in several of the theoretical approaches developing applied frameworks, is a dilution and oversimplification of concepts, that then become vulnerable to misinterpretation.

Within the ethnographic literature, there have been numerous debates about the tensions and discrepancies between the contribution ethnographers think they can make and the expectations and assumptions from the rest of the HCI community about what they ought to provide. As to their input into the design process others have commented, too, on how such fine-grained analyses of work often leads to a conservatism when it comes to considering the development and deployment of new technologies (Grudin and Grinter, 1995). Having gone to such length to reveal the richness of work practices there has been much resistance to then use these as a basis for suggesting alternative set-ups, incorporating new systems. In contrast, a trend has been to use the findings from ethnographies of the workplace to highlight the dangers of disrupting current ways of working with new technologies. For example, Heath et al (1993) discuss how existing work practices in a dealing room of the Stock Exchange would be perturbed if new technological input devices were introduced (e.g. speech recognition systems). Rogers (1992) also speculated about the problems of increasingly offloading coordination work (e.g. scheduling) of teams working together onto a computer network, based on a distributed cognition analysis of a close-knit team of engineers who had networked their PCs.

As stressed by Button and Dourish (1996) a dilemma facing researchers is that ethnomethodology's "tradition is in analyzing practice, rather than inventing the future." (p. 21). But where does this leave the ethnomethodologist or ethnographer who has moved into interaction design? To be always recorders and interpreters of events? Alternatively, is it possible for them to become more concerned with the process of design, and to shift between different levels of description, that make sense to both research and design? Hughes et al (1997) have discussed at length the communicative gap between the "expansive textual expositions of the ethnographer and the abstract graphical depictions and 'core concepts' of the designer" (p1). Button (1993) and Shapiro (1994) note, too, how the descriptive language constructed in ethnographic studies has been of little relevance to the practical problem of designing computer systems. Anderson (1994) points out how discussions about these differences can end up as sterile debates resulting in a number of misconceptions being perpetuated. These take the form whereby ethnographers are caricatured as obdurate, refusing to provide the kinds of prescriptions designers are assumed to want. The designers' needs are, conversely, caricatured as always having to be couched in a formal notation, "as if design consisted in jigsaw-puzzle solving and only certain shaped pieces were allowed" (p.153). Alternatively, Anderson has argued that a new sensibility – a fresh way of viewing design problems – is needed whereby ethnographies can provoke designers to question their current frames of reference that are currently so tied to the traditional problem-solution paradigm. In so doing, he hopes that the deadlock will be surpassed and new design possibilities ensue. In a similar vein, we saw how Button and Dourish (1996) have argued for a new synthesis to viewing design within ethnomethodology concepts and ethnomethodology concepts within technological concepts.

So how can theory best *inform* design? Are there other ways of translating theory-based knowledge, besides turning it into guidelines or analytic frameworks that end up having limited utility? It would seem that quite a different frame of reference is needed – one which focuses more on the *process* of design and how the different kinds of designers, themselves, want to be supported. In addition, a quite different perspective on the nature of the relationship between researchers and designers is needed – one which sees them working more as partners collaborating together and engaged in ongoing dialogues rather than one based on the rhetoric of compassion, where researchers are viewed as educators and purveyors of knowledge whilst designers are viewed only as recipients (Rogers, 1997). It may also be possible for researchers to become designers (and vice versa) and lead by example, facilitating knowledge transfer by being able to take both perspectives.

One way that new theoretical approaches can make more of a contribution to the practice of interaction design, therefore, is to progress further with rethinking new mechanisms of 'knowledge transfer'. As suggested earlier, the potential value of building up a lingua franca – that different parties in research and design can use to refer to common referents – is an important step in this direction. As Green *et al.* (1996) comment, "all too frequently the level of discourse in evaluating software, even between highly experienced users, is one in which important concepts are struggling for expression." (p. 105). Their hope is that the vocabulary of cognitive dimensions will offer a better means of articulating trade-offs, concerns and frustrations when designing. Utilizing poignant metaphors is another rhetorical device that could be extended for concretizing the intangible and the difficult. For example, Star's (1989) notion of 'boundary objects' to describe objects which "are plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity across sites" (p. 46) has been taken up by numerous researchers and designers as a way of better articulating previously nebulous and ill-formed ideas. Bowers and Pycock (1994) have also shown how the use of other rhetorical devices can have value for practice: outlining how the metaphorical description of 'resistances' and 'forces' can be used to express different aspects of a design space. Likewise, Rogers (1994) has used rhetorical devices together with various cognitive dimensions to analyze aspects of the design and use of groupware systems. One of the main attractions to these kinds of concept, is that they readily appropriate or map onto everyday terms and concepts that are relatively easy to understand. This allows for analogical reasoning that can be generalized across a range of topics.

Pattern languages are another form of abstraction being introduced into HCI and software engineering (Borchers, 2001; Erickson, 1999). Originally developed by the architect Christopher Alexander for describing architecture and urban design, such as aspects of both city design and wall design, they are now being taken up to describe patterns of software design and use. A major attraction of adopting these and other interconnected sets of concepts (e.g. Activity Theory) is that they provide a pro forma, for identifying abstractions that can be visualized and constructed as meaningful units of analysis.

In sum, one of the main contributions of continuing to import and develop theoretically-based approaches in HCI is as a basis from which to enable new accounts, frameworks and concepts to be constructed. In turn, these have the potential for being developed further into a more extensive design language, that can be used both in research and design. Given the increasing diversity of people now involved in the design of an increasingly diverse set of interactive products and user experiences, it would seem even more pressing for such a language(s) to be developed. This in itself, however, is no easy task. It requires determining which of the new terms, metaphors, and other abstractions are useful for articulating design concerns – and which, importantly, the different people see value in and feel comfortable using. Designers and researchers need to begin to engage in more dialogues, identifying areas of conceptual 'richness' and design 'articulation'. As part of this enterprise, the practice of interaction design, itself, would greatly benefit from further research – especially an analysis of the different languages and forms of representations that are used, together with a better understanding of the trade-offs and numerous decisions facing designers as they seek to harness the ever increasing range of technological possibilities.

Acknowledgements

This chapter is dedicated to the late Mike Scaife, whose ideas and feedback on earlier drafts were invaluable.

References

Ackerman, M. & Halverson, C. (1998). Considering an Organization's Memory. In *Proceedings of Computer-Supported Cooperative Work*, CSCW'98, ACM, New York, 39-48.

Anderson, J.R. (1983). *The Architecture of Cognition.* Harvard University Press Cambridge, MA.

Anderson, R., Button, G. and Sharrock, W. (1993). Supporting the design process within an organizational context. *Proceedings of 3rd ECSCW,* 13-17th September, Milan, Italy, Kluwer Academic Press. 47-59.

Anderson, R.J. (1994). Representations and requirements: The value of ethnography. In system design. *Human Computer Interaction,* 9, 151-182.

Atwood, M.E., Gray, W.D. & John, B.E. (1996). Project Ernestine: Analytic and Empirical methods applied to a real world CHI problem. In Rudisill, M., Lewis, C., Polson, P. and McKay, T.D, (Eds.), *Human Computer Interface Design: Success Stories, Emerging Methods and Real World Context*. Morgan Kaufmann, San Francisco. 101-121.

Bailey, B. (2000). How to improve design decisions by reducing reliance on superstition. Let's start with Miller's Magic 7+-2, *Human Factors International, Inc,* September 2000, www.humanfactors.com

Bannon, L. & Bødker, S. (1991). Encountering artifacts in use. In Carroll, J. (Ed.), *Designing Interaction: Psychology at the Human-Computer Interface,* , Cambridge University Press, New York, 27-253.

Barnard, P. (1991). Bridging between basic theories and the artifacts of Human-Computer Interaction. In Carroll, J. (Ed.), *Designing Interaction: Psychology at the Human-Computer Interface,* Cambridge University Press, New York, 103-127.

Barnard, P.J. & May, J. (1999). Representing cognitive activity in complex tasks. *Human-Computer Interaction,* 14, 93-158.

Barnard, P.J., Hammond, N., Maclean, A., & Morten, J. (1982). Learning and remembering interactive commands in a text editing task. *Behaviour and Information Technology*, 1, 347-358.

Barnard, P.J., May, J., Duke, D.J. & Duce, D.A. (2000). Systems Interactions and Macrotheory. *Transactions On Computer Human Interaction,* 7, 222-262.

Bellamy, R.K.E. (1996). Designing educational technology: computer-mediated change. In B. Nardi (Ed.) *Context and Consciousness: Activity Theory and Human-Computer Interaction.* MIT Press, Mass, 123-146.

Bellotti, V. (1988). Implications of current design practice for the use of HCI techniques. In D.M. Jones & R. Winder (Eds.) *People and Computers IV: Designing for Usability, Proc HCI'88,* CUP, 13-34.

Bentley, R., Hughes J.A., Randall, D., Rodden, T., Sawyer, P., Sommerville, I. & Shapiro, D. (1992). Ethnographically-informed systems design for air traffic control. In *Proceedings of the Conference on Computer Supported Cooperative Work, CSCW'92*, ACM, New York, 123-129.

Beyer, H. & Holtzblatt, K. (1998). *Contextual Design: Customer-Centered Systems.* Morgan Kauffman, San Francisco.

Bly, S. (1997). Field work: is it product work? *ACM Interactions Magazine,* January and February, 25-30.

Bødker, S. (1989). A human activity approach to user interfaces. *Human Computer Interaction, 4* (3), 171-195.

Borchers, J. (2001). *A Pattern Approach to Interaction Design.* Wiley, Chichester.

Bowers, J. & Pycock, J. (1994). Talking through design: requirements and resistance in cooperative prototyping. In *CHI'94 Conference Proceedings,* ACM, New York, 299-305.

Button, G. & Dourish, P. 1996. Technomethodology: Paradoxes and possibilities. In *CHI'96 Conference Proceedings,* ACM, New York, 19-26.

Button, G. (1993). (Ed.). *Technology in Working Order.* Routledge, London.

Button, G. (1997). Book review: Cognition in the Wild, *CSCW,* 6, 391-395.

Card, S. K., Moran, T.P. & Newell, A. (1983). *The Psychology of Human-Computer Interaction.* Hillsdale, LEA, New Jersey.

Card, S.K., Mackinlay. J.D. & Shneiderman, B. (1999). Information Visualization. In Card, S.K., Mackinlay. J.D. & Shneiderman, B. (Eds.), *Readings in Information Visualization*. Morgan Kaufman Publishers, SF, USA, 1-35.

Carroll, J.M. (1991). (Ed.) *Designing Interaction: Psychology at the Human-Computer Interface.* Cambridge University Press, Cambridge.

Carroll, J.M., Kellogg, W.A. & Rosson, M.B. (1991). The Task-Artifact Cycle. In J. Carroll (Ed.) *Designing Interaction: Psychology at the Human-Computer Interface.* Cambridge University Press, Cambridge. 74-102.

Castell, F. (2002) Theory, theory on the wall…. *CACM,* 45, 25-26. *Cognitive Science,* 18, 87-122.

Clancey, W.J. (1997). *Situated Cognition: On Human Knowledge and Computer Representations.* Cambridge University Press, Cambridge.

Clemmensen, T. & Leisner, P. (2002), Community Knowledge in an emerging online professional community: the interest of theory among Danish usability professionals. *IRIS'25 (International Systems Research in Scandinavia).* 10-13th August, Kulhuse, Denmark.

Collins, P. Shulka, S., & Redmiles, D. (2002). Activity theory and system design: a view from the trenches. *CSCW,* 11, 55-80.

Cooper, G. (1991). *Representing the User.* Unpublished PhD, Open University, UK.

Draper, S. (1992). Book review: The New Direction for HCI? "Through the Interface: A Human Activity Approach to User Interface Design," by S. Bodker. *International Journal of Man-Machine Studies,* 37(6), 812-821.

Engestrøm, Y. & Middleton, D. (1996). (Eds.) *Cognition and Communication at Work.* Cambridge University Press, Cambridge, UK.

Engestrøm, Y. (1990). *Learning, Working and Imagining: Twelve Studies in Activity Theory*. Orienta-Konsultit, Helsinki.

Engestrøm, Y. (1993). Developmental studies of work as a test bench of activity theory: the case of primary care medical practice. In *Understanding Practice: Perspectives on Activity and Context.* S. Chaiklin and J. Lave, Eds., Cambridge University Press, Cambridge, UK, 64-103.

Erickson, T. (1999). Towards a Pattern Language for Interaction Design. In P. Luff, J. Hindmarsh, & C. Heath (Eds.) *Workplace Studies: Recovering Work Practice and Informing Systems Design*. Cambridge University Press, Cambridge. 252-261.

Erickson, T. (2002). Theory Theory: A Designer's View, *CSCW,* 11, 269-270.

Fitzpatrick, G. (2003) *The Locales Framework: Understanding and Designing for Wicked Problems.* Kluwer, The Netherlands.

Fjeld, M., Lauche, K., Bichsel, Voorhorst, F., Krueger, H. & Rauterberg, M. (2002). Physical and virtual tools: activity theory applied to the design of groupware. *CSCW*, 153-180.

Flor, N.V. & Hutchins, E. 1992. Analyzing distributed cognition in software teams: a case study of collaborative programming during adaptive software maintenance. In J. Koenemann-Belliveau, T. Moher, & T. Robertson, (Eds.) *Empirical Studies of Programmers: Fourth Workshop,* Ablex, Norwood, NJ, 36-64.

Gabrielli, S., Rogers, Y. & Scaife, M. (2000). Young Children's Spatial Representations Developed through Exploration of a Desktop Virtual Reality Scene, *Education and Information Technologies,* 5(4), 251-262.

Garbis, C. & Waern, Y. (1999). Team Co-ordination and Communication in a Rescue Command Staff - The Role of Public Representations, *Le Travail Humain,* 62 (3), Special issue on Human-Machine Co-operation, 273-291.

Garfinkel, H. & Sacks, H. (1970). On the formal structures of practical action. In J. McKinney and E. Tiryakian, Eds., Appleton-Century-Crifts *Theoretical Sociology*, New York, 338-386.

Garfinkel, H. (1967). *Studies in Ethnomethodology.* Polity Press, Cambridge.

Gaver, B. (1991). Technology affordances. In *CHI'91 Conference Proceedings,* 85-90, Addison-Wesley, Reading MA.

Geertz, C. (1993). *The Interpretation of Cultures: Selected Essays.* Fontana Press, London.

Gibson, J.J. (1966). *The Senses Considered as Perceptual Systems.* Houghton-Mifflin, Boston.

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

Green, T. R. G. (1990). The cognitive dimension of viscosity: a sticky problem for HCI. In D. Diaper, D. Gilmore, G. Cockton & B. Shakel, (Eds.), *Human-Computer Interaction - INTERACT'90.* Elsevier Publishers, B.V., North Holland. 79-86.

Green, T.R.G. (1989). Cognitive dimensions of notations. In A. Sutcliffe & L. Macaulay (Eds.), *People and Computers V.*, Cambridge University Press, Cambridge, 443-459.

Green, T.R.G., Davies, S.P. & Gilmore, D.J. (1996). Delivering cognitive psychology to HCI: the problems of common language and of knowledge transfer, *Interacting with Computers*, 8 (1), 89-111.

Grudin, J. & Grinter, R.E. (1995). Commentary: Ethnography and Design*, CSCW,* 3, 55-59.

Grudin, J. (2002). HCI theory is like the public library. Posting to CHIplace online discussion forum, Oct 15th 2002, *www.chiplace.org*

Gunther, V.A., Burns, D.J. and Payne, D.J. (1986). Text editing performance as a function of training with command terms of differing lengths and frequencies. *SIGCHI Bulletin,* 18, 57-59.

Halloran, J., Rogers, Y. and Scaife, M. (2002). Taking the 'No' out of Lotus Notes: Activity Theory, Groupware and student work projects, In *Proc. of CSCL,* Lawrence Erlbaum Associates, Inc. NJ.,169-178.

Halverson, C.A. (1995). Inside the cognitive workplace: new technology and air traffic control. *PhD Thesis*, Dept. of Cognitive Science, University of California, San Diego, USA.

Halverson, C.A. (2002). Activity theory and distributed cognition: Or what does CSCW need to DO with theories? *CSCW,* 11, 243-275.

Heath, C. & Luff, P. (1991). Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. In *Proceedings of the Second European Conference on Computer-Supported Cooperative Work*, Kluwer, Dordrecht, 65-80.

Heath, C. Jirotka, M., Luff, P. & Hindmarsh, J. (1993). Unpacking collaboration: the international organisation of trading in a city dealing room. In *Proceedings of the*

Third European Conference on Computer-Supported Cooperative Work, Kluwer, Dordrecht, 155-170.

Hollan, J., Hutchins, E. & Kirsh, D. (2000). Distributed Cognition: toward a new foundation for human-computer interaction research. *Transactions on Human-Computer Interaction,* 7(2), 174-196.

Hughes, J.A., O'Brien, J., Rodden, T. & Rouncefield, M. (1997). CSCW and Ethnography: A presentation framework for design. In I. McClelland. G. Olson, G. van der Veer, A. Henderson & S. Coles, (Eds.) *Proceedings of the conference on Designing Interactive Systems: Processes, Practices and Techniques (DIS'97, Amsterdam, The Netherlands, Aug 18 - 20),* ACM Press, New York, 147 -158.

Hutchins, E. & Klausen, T. (1996). Distributed Cognition in an Airline Cockpit. In D. Middleton, & Y. Engeström (Eds.), *Communication and Cognition at Work*, Cambridge University Press, Cambridge. 15-34.

Hutchins, E. & Palen, L. (1997). Constructing Meaning from Space, Gesture and Speech. In L.B. Resnick, R. Saljo, C. Pontecorvo, & B. Burge, (Eds.) *Discourse, Tools, and Reasoning: Essays on Situated Cognition.* Springer-Verlag, Heidelberg, Germany. 23-40.

Hutchins, E. (1995). *Cognition in the Wild.* MIT Press, Mass.

Hutchins, E., Hollan, J.D. and Norman, D. (1986). Direct manipulation interfaces. In S. Draper and D. Norman, (Eds.) *User Centred System Design*. Lawrence Erlbaum Associates, NJ., 87-124.

Kaptelinin, V., Nardi, B.A. and Macaulay, C. (1999). The Activity checklist: a tool for representing the "space" of context. *Interactions* july+august 1999, 27-39.

Kaptelinin, V. (1996). Computer-mediated activity: functional organs in social and developmental contexts. In B. Nardi (Ed.), *Context and Consciousness: Activity Theory and Human-Computer Interaction*, MIT Press, Mass. 45-68.

Kieras, D. & Meyer, D.E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction,* 12, 391-438.

Kieras, D. & Polson, P.G. (1985). An approach to the formal analysis of user complexity. *International Journal of Man-Machine Studies,* 22, 365-394.

Kieras, D. (1988). Towards a practical GOMS model methodology for user-interface design. In M. Helander (ed.) *Handbook of Human-Computer Interaction*. Amsterdam: North-Holland, 135-157.

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

Kirsh, D. (1997). Interactivity and Multimedia Interfaces. *Instructional Science*, 25, 79-96.

Kirsh, D. (2001). The context of work. *HCI*. 6(2), 306-322.

Kuutti, K. (1996). Activity Theory as a potential framework for Human-Computer Interaction research. In B. Nardi (Ed.), *Context and Consciousness: Activity Theory and Human-Computer Interaction*, MIT Press, Mass. 17-44.

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. Carroll (Ed.), *Designing Interaction: Psychology at the Human-Computer Interface.* Cambridge University Press, New York, 60-73.

Ledgard, H., Singer, A., & Whiteside, J. (1981). Directions in human factors for interactive systems. In G. Goos & J. Hartmanis (Eds.) *Lecture Notes in Computer Science,* 103, Berlin, Springer-Verlag.

Leontiev, A.N. (1978). *Activity, Consciousness and Personality*. Prentice Hall.

Lewis, C., Polson, P., Wharton, C., and Rieman, J. (1990). Testing a walkthrough methodology for theory-based design of walk-up-and-use interfaces. *CHI'01 Proceedings,* New York: ACM, 137-144.

Long, J. & Dowell, J. (1989). Conceptions for the discipline of HCI: craft, applied science, and engineering. In A. Sutcliffe & L. Macaulay (Eds). *People and Computers V*, CUP, Cambridge, UK, 9-32.

Long, J. & Dowell, J. (1996). Cognitive engineering human-computer interactions. In *The Psychologist*, July, 313-317.

Mantovani, G. (1996). Social Context in HCI: A new framework for mental models, cooperation and communication. *Cognitive Science,* 20, 237-269.

Masterman, E. & Rogers, Y. (2002). A framework for designing interactive multimedia to scaffold young children's understanding of historical chronology. *Instructional Science*. 30, 221-241

Microsoft (2002). Consumer Input, Scientific Analysis Provide Foundation for MSN 8 Research and Innovation. Redmond, Wash., Oct 23rd, 2002, *Presspass, www.microsoft.com,* p.4.

Modugno, F.M., Green, T.R.G. & Myers, B. (1994). Visual Programming in a visual domain: a case study of cognitive dimensions. In G. Cockton, S.W. Draper, & G.R.S. Weir, (Eds.) *People and Computers IX.* Cambridge University Press, Cambridge, UK.

Mohlich, R. and Nielsen, J. (1990). Improving a human-computer dialogue. *Comm. of the ACM,* 33(3), 338-48.

Monk, A. (1984) (Ed.). *Fundamentals of Human-Computer Interaction*. Academic Press, London.

Nardi, B.A. & Johnson, J. (1994). User preferences for task-specific versus generic application software. In *CHI'94 Proc.* ACM, New York, 392-398.

Nardi, B.A. (1996). (Ed.) *Context and Consciousness: Activity Theory and Human-Computer Interaction.* MIT Press, Mass.

Nardi, B.A. (2002). Coda and response to Christine Halverson. *CSCW,* 269-275.

Neisser, U. (1985). Toward an ecologically oriented cognitive science. In T.M. Schlecter, & M.P. Toglia (Eds.) *New Directions in Cognitive Science,* Ablex Publishing Corp, Norwood, N.J. 17-32.

Norman, D. (1986). Cognitive engineering. In S. Draper and D. Norman, (Eds.) *User Centered System Design*. Lawrence Erlbaum Associates, NJ, 31-61.

Norman, D. (1988). *The Psychology of Everyday Things.* Basic Books, NY.

Norman, D. (1993). Cognition in the head and in the world. *Cognitive Science,* 17 (1), 1-6.

Norman, D. (1999). Affordances, conventions and design. *Interactions,* may/june 1999, 38-42. ACM, New York.

Oliver, M. (1997). *Visualisation and manipulation tools for Modal logic*. Unpublished PhD thesis, Open University.

Olson, J.S. & Moran, T.P. (1996). Mapping the method muddle: guidance in using methods for user interface design. In M. Rudisill, C. Lewis, P. Polson, & T.D McKay, (Eds.) *Human Computer Interface Design: Success Stories, Emerging Methods and Real World Context*. Morgan Kaufman, San Francisco 269-300.

Olson, J.S. & Olson, G.M. (1991). The growth of cognitive modeling since GOMS. *Human Computer Interaction*, 5, 221-266.

Otero, N (2003). Interactivity in graphical representations: assessing its benefits for learning. *Unpublished Dphil,* University of Sussex, UK.

Patel, V. L. & Groen, G.J. (1993). Comparing apples and oranges: some dangers in confusing frameworks and theories. *Cognitive Science,* 17, 135-141.

Pirolli, P. & Card, S. (1997). The evolutionary ecology of information foraging. *Technical report, UIR-R97-01.* Palo Alto Research Center, CA.

Plowman, L., Rogers, Y. & Ramage, M. (1995). What are workplace studies for? In *Proc of the Fourth European Conference on Computer-Supported Cooperative Work.* Dordrecht, The Netherlands, Kluwer., 309-324.

Polson, P.G., Lewis, C., Rieman, J. & Wharton, C. (1992). Cognitive walkthroughs: a method for theory-based evaluation of user interfaces. *International Journal of Man-Machine Studies,* 36, 741-73.

Preece, J., Rogers, Y., Sharp, H., Benyon, D. Holland, S. & Carey, T. (1994). *Human-Computer Interaction*. Addison-Wesley, London.

Price, S. (2002). Diagram representation: the cognitive basis for understanding animation in education. *Unpublished Dphil,* University of Sussex, UK.

Rasmussen, J. & Rouse, W. (1981) (Eds.). *Human Detection and Diagnosis of System Failures.* Plenum Press, New York.

Rasmussen, J. (1986). *On Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering.* Elsevier, Amsterdam.

Rodden, T., Rogers, Y., Halloran, J. & Taylor, I. (2003). Designing novel interactional work spaces to support face to face consultations. To appear in *CHI Proc.,* ACM.

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

Rogers, Y. (1992). Ghosts in the network: distributed troubleshooting in a shared working environment. In *CSCW'92 Proc.,* ACM, New York, 346-355.

Rogers, Y. (1993). Coordinating computer mediated work. *CSCW,* 1, 295-315.

Rogers, Y. (1994). Exploring obstacles: Integrating CSCW in evolving organisations. In *CSCW'94 Proc.,* ACM, New York, 67-78.

Rogers, Y. (1997). Reconfiguring the Social Scientist: Shifting from telling designers what to do to getting more involved. In G.C. Bowker, S.L. Star, W. Turner & L. Gasser. (Eds.)*, Social Science, Technical Systems and Cooperative Work*, LEA, 57-77.

Rogers, Y., Bannon, L. & Button, G. (1993). Rethinking theoretical frameworks for HCI. *SIGCHI Bulletin,* 26(1), 28-30.

Rogers, Y. and Scaife, M. (1998). *How can interactive multimedia facilitate learning?* In Lee, J. (ed.) Intelligence and Multimodality in Multimedia Interfaces: Research and Applications. AAAI. Press: Menlo Park, CA.

Rogers, Y., Preece, J. & Sharp, H. (2002). *Interaction Design: Beyond Human-Computer Interaction.* Wiley, New York.

Scaife, M. & Rogers, Y. (1996). External Cognition: how do graphical representations work*? International Journal of Human-Computer Studies,* 45, 185-213.

Scaife, M., Halloran, J., and Rogers, Y.(2002) Let's work together: supporting two-party collaborations with new forms of shared interactive representations. *Proceedings of COOP'2002.* Nice, France, August 2002, IOS Press, The Netherlands,123-138.

Scapin, D.L. (1981). Computer commands in restricted natural language: some aspects of memory of experience. *Human Factors,* 23, 365-375.

Shapiro, D. (1994). The limits of ethnography: combining social sciences for CSCW. In *Proc of CSCW'94,* ACM, NY. 417-428.

Shneiderman, B. (1992) *Designing the User Interface: Strategies for Effective Human-Computer Interaction. 2nd Ed.,* Reading, MA: Addison-Wesley.

Shneiderman, B. (2002a). *Leonardo's Laptop.* MIT Press.

Shneiderman, B. (2002b). HCI theory is like the public library. Posting to CHIplace online discussion forum, Oct 15th 2002, *www.chiplace.org*

Spasser, M. (2002). Realist Activity Theory for digital library evaluation: conceptual framework and case study. *CSCW,* 11, 81-110.

St. Amant, R. (1999). User Interface affordance in a planning representation. *Human-Computer Interaction*, 14, 317-354.

Star, S.L. (1989). The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving. In L. Gasser, L. & M.N. Huhns (Eds.) *Distributed Artificial Intelligence, Volume II,* Morgan Kaufmann, SF Mateo, CA. 37-54.

Star, S.L. (1996). Working together: Symbolic interactionism, activity theory and information systems. In Y. Engestrøm, & D. Middleton, (Eds.) *Cognition and Communication at Work.* Cambridge, UK, CUP. 296-318.

Suchman, L.A. (1983). Office Procedure as Practical Action: Models of Work and System Design. *TOIS,* 1(4), 320-328.

Suchman, L.A. (1987). *Plans and Situated Actions.* Cambridge University Press, Cambridge.

Sutcliffe, A. (2000). On the effective use and reuse of HCI knowledge. *Transactions on Computer-Human Interaction,* 7 (2), 197-221.

Vera, A.H. & Simon, H.A. (1993). Situated Action: A Symbolic Interpretation. *Cognitive Science,* 17 (1), 7-48.

Vicente, K.J. & Rasmussen, J. (1990). The ecology of man-machine systems II: Mediating 'direct perception' in complex work domains. *Ecological Psychology*, 2, 207-249.

Vicente, K.J. (1995). A few implications of an ecological approach to human factors. In J. Flach, P. Hancock, J. Carid, & K.J. Vicente, (Eds.) *Global Perspective on the Ecology of Human-Machine Systems,* 54-67.

Weiser, M. (1991). The Computer for the 21st Century. *Scientific American,* 265 (3), 94-104.

Winograd, T. (1997). From Computing Machinery to Interaction Design. In P. Denning and R. Metcalfe (Eds.) *Beyond Calculation: the next fifty years of computing.* Springer-Verlag, 149-162.

Wood, C. A (1995). Cultural-Cognitive Approach to Cognitive Writing. *Unpublished DPhil dissertation,* University of Sussex, UK.

Woods, D.D. (1995). Toward a theoretical base for representation design in the computer medium: ecological perception and aiding cognition. In J. Flach, P. Hancock, J. Carid, & K.J. Vicente, (Eds.) *Global Perspective on the Ecology of Human-Machine Systems,* 157-188.

Wright, P., Fields, R. & Harrison, M. (2000). Analyzing Human-Computer Interaction As Distributed Cognition: The Resources Model. *Human Computer Interaction,* 51(1), 1-41.

Yang, S. Burnett, M.M., Dekoven, E. & Zloof, M. (1995). Representations design benchmarks: a design-time aid for VPL navigable static representations. *Dept. of Computer Science Technical Report 95-60-4* Oregon State University (Corvallis).

Zhang, J. & Norman, D.A. (1994). Representations in distributed cognitive tasks. *Cognitive Science,* 18, 87-122.