

Article

Putting ethnography to work: the case for a cognitive ethnography of design

Ball, Linden and Ormerod, Thomas C.

Available at <http://clock.uclan.ac.uk/6330/>

Ball, Linden ORCID: 0000-0002-5099-0124 and Ormerod, Thomas C. (2000) Putting ethnography to work: the case for a cognitive ethnography of design. International Journal of Human-Computer Studies, 53 (1). pp. 147-168. ISSN 1071-5819

It is advisable to refer to the publisher's version if you intend to cite from the work.

<http://dx.doi.org/10.1006/ijhc.2000.0372>

For more information about UCLan's research in this area go to <http://www.uclan.ac.uk/researchgroups/> and search for <name of research Group>.

For information about Research generally at UCLan please go to <http://www.uclan.ac.uk/research/>

All outputs in CLoK are protected by Intellectual Property Rights law, including Copyright law. Copyright, IPR and Moral Rights for the works on this site are retained by the individual authors and/or other copyright owners. Terms and conditions for use of this material are defined in the [policies](#) page.

Running Head: Putting ethnography to work

Putting ethnography to work: The case for a cognitive ethnography of design

Linden J. Ball

University of Derby, UK

Thomas C. Ormerod

Lancaster University, UK

Ball, L.J., & Ormerod, T.C. (2000). Putting ethnography to work: The case for a cognitive ethnography of design. *International Journal of Human-Computer Studies*, 53, 147-168.

Correspondence to: Linden J. Ball, Institute of Behavioural Sciences, University of Derby, Mickleover, Derby, DE3 5GX. Email: L.J.Ball@derby.ac.uk

Putting ethnography to work: The case for a cognitive ethnography of design

Linden J. Ball

Institute of Behavioural Sciences, University of Derby
Mickleover, Derby DE3 5GX, UK
L.J.Ball@derby.ac.uk

Thomas C. Ormerod

Department of Psychology, Lancaster University,
Fylde College, Lancaster LA1 4YF UK
T.Ormerod@lancaster.ac.uk

Abstract

The methods of ethnography and cognitive psychology are frequently set in opposition to each other. Whilst such a view may be appropriate in defining pure, or prototypical, classes of each activity, the value and necessity of such a distinction is broken down when researchers are goal-directed to study complex work domains in order to foster technological change. In this paper, we outline a rapprochement of these methods, which we term cognitive ethnography. The value of qualifying ethnography in this way is to emphasise systematically the differences between ethnography as a radial category and the kinds of legitimate method used to study work practices which are often referred to as ethnographic, but which in practice differ in important ways from prototypical ethnographic studies. Features of cognitive ethnography such as observational specificity, verifiability and purposiveness challenge many of the tenets of a pure ethnographic method, yet they are essential for studies that are undertaken to inform technological change. We illustrate our arguments with reference to a project to develop a tool for supporting design re-use in innovative design environments.

Since the seminal work of Simon (1969), the design of technological artefacts has been increasingly characterised as a cognitive activity, and many researchers have employed the observational and experimental methods of cognitive psychology in understanding the nature of design expertise and in developing and evaluating tools for supporting the design process (e.g., Guindon, 1990; Visser, 1990). There are a number of reasons why design might be seen as the territory of cognitive psychologists, of which two concern us here. First, as Simon (1969) points out, design represents the archetypal complex problem-solving activity undertaken by individuals. In this sense, cognitive psychologists use design as a domain in which to further their understanding of 'cognition in-the-head'. Second, the hypothetico-deductive research tradition that underpins cognitive psychology can be seen as providing an analogy for design practice. Carroll and Campbell (1989) describe the process of designing technological devices as a task-artefact cycle, in which the analysis of tasks motivates technological developments, which in turn affect the nature and conditions of task performance. In essence, artefacts are implicit psychological hypotheses that are tested through subsequent empirical evaluation.

Whilst it may appear natural to conceptualise design as a cognitive activity, design is also typically a highly context-bound endeavour, situated, as it is, within commercial organisations whose practices, structures and social interactions shape both final design products and the processes by which those products are achieved (cf. Buchiarelli, 1988). Given the highly contextualised nature of design activity, it remains paradoxical that the majority of existing studies of design expertise have ignored the role of situational and social factors in design in preference to carrying out laboratory-style investigations in which such factors are controlled for. Manipulations instead tend to focus on variables such as the designer's skill or experience level or the type of problem that is being tackled (e.g., routine or non-routine). This rather narrow, de-contextualised focus of much design research is also echoed in much of the traditional research in other domains of human skill such as medical diagnosis. As Hoffman, Feltovich and Ford (1997) note, however, research is beginning to reveal that constraints and opportunities

arising from contextual and social factors are at least as important determinants of the exercise of expertise as are task-specific and skill-specific variables. Thus there is a need for methodologies beyond laboratory experimentation.

A context-sensitive technique with which cognitive psychologists are familiar is verbal protocol analysis. Ericsson and Simon (1993) claim that verbal report data can be collected in situ without necessarily interfering with task performance, to offer a complete, accurate and objective account for a wide range of cognitive activities, notably the kinds of problem-solving activity thought to underlie design (e.g., Simon, 1969; Ball & Ormerod, 1995). Whilst these arguments may be open to dispute (e.g., Nisbett & Wilson, 1977), there is no doubt that verbal protocol analysis has proved a valuable method for the exploration of cognitive performance in both pure and applied task domains. It remains the case, however, that such methods are typically used to understand single-person cognition in socially-impoverished environments, rather than multi-agent cognition in full-blown, people-rich environments. Problem solving which is contextualised by other people (e.g., problem solving which arises in team-oriented environments such as hospitals, businesses and design groups) would seem to require methods of analysis that are often unfamiliar coinage for cognitive psychologists.

Requirements for observational analyses of situated behaviours suggest that a technique such as ethnography (e.g., Hammersley & Atkinson, 1983), which has been used in sociological and anthropological research, might usefully be employed. Ethnography aims to provide accounts of activity as perceived and recognised by the individuals present within the situation as participant observers. This approach affords interesting advantages over traditional cognitivist techniques in that it is geared toward gleaning insights on important aspects of human activity which more scientific techniques tend to miss -- including behaviours that arise at the social interface or behaviours which may not be directed towards the attainment of organisational or managerial goals.

A key question faced by design researchers and practitioners, as well as those from other areas of human factors, is the extent to which ethnography can, or indeed

should, be adapted so as to render it sensitive to the aims of cognitive psychology. There are two key areas in which there is an apparent clash between the aims of ethnography and cognition, which map onto the reasons espoused above for the cognitive ‘territorialisation’ of design. First, the goal of understanding cognition in-the-head (i.e., processes, mechanisms, structures and representations that are inherently mental in nature) seems diametrically opposed to the study of cognition as a disembodied cultural construct, which relates more closely to the goals of sociology or anthropology, the disciplines associated with ethnographic methods. Second, the analysis-for-a-purpose goals that characterise applied cognitive psychology are at odds with an apparent requirement for ethnography to be as theoretically and pragmatically neutral as possible.

In the present paper we suggest how ethnographic methods might be used successfully in investigating the interplay between people-laden contexts and expert cognition, whilst servicing the basic theoretical and applied goals of cognitive psychology. We first consider the nature of ethnography as a research methodology, and its possible uses in research on design cognition. In discussing the prospects and limitations of traditional ethnography in the analysis of design, we argue that significant problems reside in its use for goal-directed applied research. The disadvantages of ‘genuine’ ethnography lead us to propose a much more restricted form of cognitive ethnography for human factors research. A key aim for our paper is to explore the differences between pure ethnography and the qualified form of cognitive ethnography espoused here. We illustrate the value of a cognitive ethnographic approach with reference to research that we are currently undertaking, which is aimed at supporting information re-use in innovative design contexts through intelligent, computer-based support.

Ethnography as a field research method

The ways in which researchers define ethnography are varied and fuzzy. Some equate ethnography with any method that involves an element of in situ participant observation (e.g., Taylor, 1994). For others, the crux of ethnography is not so much its basis in participant observation but its epistemological stance, which is firmly located within a

socio-cultural frame of reference (e.g., Hutchins, 1995). Although each definition of ethnography is interesting in its own right, no single definition that we have encountered denotes the full set of necessary and sufficient attributes that define this concept.

Such difficulties in defining the concept of ethnography lead us to propose that it is better understood as a radial category which exhibits prototype effects (see Lakoff, 1987). Radial categories encompass a central, prototypical case as well as a number of variations from this prototype. In this way, the prototypical case of an ethnographic method can be viewed as a template against which to compare other, possibly less typical, instances of ethnography. To the extent that any observational method varies from this prototype (e.g., by omission of one or more attributes) then it is a less central case of the ethnographic technique. Of course, under this conceptualisation, issues remain as to how far one can deviate from a central case in order for a method still to be referred to as ethnographic in nature. This is an important concern, which has led to our view that at some point of deviation from a central case one has to recognise that a method may no longer be legitimately referred to as ethnography but must be referred to by a different term that acknowledges both its relationship to the central case and its differences to this prototype.

Our examination of a range of contemporary overviews of ethnography (e.g., Taylor, 1994; Toren, 1996; Rachel, 1996; Baszanger & Dodier, 1997) leads us to suggest that a prototypical case of ethnography would be characterised by the following ten features (the order in which these are listed is arbitrary):

1. Situatedness - Data are collected by a participant observer who is located within the everyday context of interest (e.g., a community of practitioners).
2. Richness - The observer studies behaviour in all manifestations, such that data are gathered from a wide range of sources including interviews, team discussions, incidental conversations, documents, as well as non-verbal interactions.
3. Participant Autonomy - The observees are not required to comply in any rigid, pre-determined study arrangements.

4. Openness - The observer remains open to the discovery of novel or unexpected issues that may come to light as a study progresses.
5. Personalisation - The observer makes a note of their own feelings in relation to situations encountered during data collection and analysis.
6. Reflexivity - The observer takes a reflective and empathetic stance in striving toward an understanding of the observee's point of view, the observer taking account of, rather than striving to eliminate, their own affects upon the behaviour of the observees.
7. Self-reflection - The observer must acknowledge that any interpretative act is influenced by the tradition to which they themselves belong.
8. Intensity - Observations must be intensive and long-term, such that the observer should become immersed in the ongoing culture of the observee's environment.
9. Independence - The observer must not be constrained by pre-determined goal-set, mind-set or theory.
10. Historicism - The observer aims to connect observations to a backdrop of historical and cultural contingencies.

These ten characteristics of a prototypical case of ethnography can be used to identify whether a method that has been employed by a researcher is a more or less central form of the ethnographic technique. For example, Woolgar's (1995) company-based study of the processes underpinning the design and production of a new range of microcomputers fits this template very well. Woolgar joined this design project for an eighteen-month period as a managerial assistant with responsibility for liaison between different sections of the company. As such, his status was that of a genuine participant observer with legitimate access across disparate sections of the company to rich, varied and situated data, including those deriving from usability trials. Woolgar's ethnography indicates a highly open stance to data acquisition, as exemplified by the multiplicity of perspectives sought on the nature of potential end-users of the microcomputer range and the weight that is attributed to unique and unexpected comments and viewpoints. The narrative associated with Woolgar's 'ethnography of computers' also reveals a reflexive

and personalised tone with its emphasis on the view that microcomputer design can be understood as a process of configuring intended users, where the notion of configuring subsumes processes such as setting constraints upon users' likely future actions.

Woolgar's study, then, is very much a prototypical ethnography which reflects the benefits which derive from a prototypical ethnography such as richness of the narrative, socio-cultural grounding, and epistemological relativism.

The features of cognitive ethnography

There is no doubt that the application of a prototypical ethnographic method and the 'thick description' (Rouncefield, Hughes & O'Brien) that it affords, enables insights that would undoubtedly be missed by methods lacking some of the features outlined above.

One cannot easily conceive how Woolgar's (1995) reanalysis of the concept of 'users' would have emerged in the absence of any of the features of a prototypical ethnography.

However, we are concerned that ethnographic criteria such as our ten-point feature list are unlikely to be applied successfully where human factors researchers are goal-directed to inform the development of technological artefacts.

There are three reasons for our concern: First, the intensity of ethnographic data-acquisition, which typically revolves around highly resource-intensive participant observation (some ethnographies spanning several years), may not be cost-effective or even possible for most projects. Pursuing applied goals such as supporting, assisting or training aspects of expert cognition typically means that there is a need to acquire an understanding of skilled behaviours within a relatively short time frame in order to effect facilitatory interventions as quickly as possible. Second, the general lack of verifiability of data analysis, which centres around story-telling in relation to people's historical and cultural endowments, presents a major problem. Human factors research (and the systems development it typically leads to) is expensive, and sponsors are quite reasonably going to seek reassurance that the design suggestions emerging from empirical studies are strongly grounded. Also, commercial systems development projects are typically pursuing generalisable technological solutions, such that findings from empirical research must be validated across a number of contexts. Third, the

independence of ethnographic data necessitates analyses based around gross descriptions of collective behaviours in terms of themes of interest rather than the testing of specific research hypotheses or theoretical concepts. Yet the specification of a technological system, even at the earliest stages of needs analysis, is (if only tacitly) setting up specific hypotheses that the researcher must test.

Consequently, we refer to the kinds of method used by ourselves as cognitive ethnography, which is, in essence, the adoption of a subset of the features of an ethnographic approach, and the deliberate violation of other features, in the pursuit of cognitive research goals. As well as there being a definable research tradition of cognitive psychology (and latterly, cognitive science), ‘cognitive’ has been used elsewhere as a qualifier. For example, the term cognitive anthropology is used to define a discipline of cross-cultural studies of conceptual categories and other cognitive phenomena (e.g., D'Andrade, 1990). To attach the qualifier ‘cognitive’ to a method rather than a discipline is less common, and is, as we shall see in the next section, deliberately challenging to generally accepted views of both cognitive and ethnographic methods. It is interesting to note that Hutchins (1995), in the last pages of his monograph (p. 371), has dubbed his own approach of describing cognition in culturally constituted settings as ‘cognitive ethnography’. However, much remains to be done in terms of defining what constitutes a cognitive ethnographic approach to observation, and how it might be used in systems design.

Our own conception of cognitive ethnography is characterised by three key features. First, it relies on small-scale data collection based around representative time slices of situated activity. As such, it demonstrates observational specificity, as opposed to the intensity of a prototypical ethnography. Second, it is purposive, in that its mode of questioning focuses on issues that are informed by some intention to intervene with, or somehow affect, existing work practices. In some instances, such as the case study that we describe later in the paper, these intentions are inherently cognitive in nature, such as understanding the interaction between expert designers’ problem solving strategies and external information repositories (including other designers). Thus, cognitive

ethnography violates the feature of independence (though not, at least in principle, the feature of openness). Third, it places a strong emphasis on verifiability, in terms of validating observations across observers, data sets and methodologies. As such, it challenges the prototypical ethnographic emphasis on personalisation, and at very least requires the explication of the role played by reflexivity in data capture and analysis. Cognitive ethnography is especially well suited to addressing applied questions of the type which dominate research on human factors, since it maintains levels of objectivity that enable replication by other observers, as well as the validation of its findings through a process of methodological triangulation involving experimental methods.

A critic might argue that the phrase ‘cognitive ethnography’ has an oxymoronic quality. Many of the features of a pure ethnography are in flagrant violation of basic tenets of the dominant paradigm of cognitive experimentation. Notably, one might see the ethnographic features of situatedness, participant autonomy, personalisation, reflexivity and independence as implying precisely the converse of a list of features that define the methods of experimental cognitive psychology, namely reductionism, experimental manipulation, impersonalisation, objectivity, and hypothesis-driven data sampling. However, we have already argued that the methods of experimental cognitive psychology are not always appropriate for exploring contextually rich multi-agent domains such as design, and that the methods of pure ethnography lack feasibility in the human factors domain. So to line up the extreme cases against each other in this way makes little sense. The proof of the method of cognitive ethnography, and probably the only sensible response to this linguistically plausible but otherwise vacuous critique, lies in the products of an application of the method.

Dimensions of ethnographies

The foregoing description of prototypical and cognitive ethnographies presents a space in which different forms of ethnographic studies can be located. However, this is not a straightforward process, for two reasons. First, there is no precise parameterisation with which ethnographic approaches can be compared; violations of one feature may be more important than violations of another. Second, there is no single continuum between

prototypical and cognitive ethnographies. There are likely to be many high-level dimensions on which ethnographies differ, of which two, 'goal' and 'focus', are of greatest interest to us. It is the placement within these dimensions, we suggest, that is the principal determinant of whether a study embodies features of a prototypical ethnography.

The first dimension ranges from 'theoretical' to 'applied', and aims to capture an essential aspect of the research goals motivating the ethnographic study. Woolgar's research, for example, is primarily directed toward gleaning fundamental theoretical insights into microcomputer design and production processes rather than an understanding of value for applied work (e.g., supporting the design process through new tools or methods). Likewise the seminal work on situated action and distributed cognition (e.g., Lave, 1988; Lave, Murtaugh & de la Rocha, 1984; Suchman, 1987; Goodwin & Goodwin, 1995; Hutchins, 1995) also appears to have been primarily motivated by the pursuit of theoretical rather than applied goals. Such studies tend to map closely onto the ten features we identify as characterising the radial category of prototypical ethnography. By the same token, none of these studies reflects strongly the features of purposiveness, specificity and verifiability that we argue characterise cognitive ethnography. It might be argued, however, that, to some extent, all studies have these latter features. For example, Hutchins' (1995) studies of navigation were partly supported by the US Office of Naval Research who presumably had some interest in using the results of his work - thus it had some aspects of purpose. On closer inspection, however, his work appears to have been mainly guided by a broad requirement to study the richness underlying navigation as a work domain, and not by a pre-planned intervention within that work domain.

Some contemporary researchers have employed more restricted forms of ethnographic techniques in their field research in pursuit of systems design. For example, Hughes and colleagues (e.g., Hughes, King, Rodden & Anderson, 1994) have referred to 'quick and dirty' or 'lightweight' ethnography as being invaluable in speeding up the delivery of ethnographic results. They argue that the slow pace of purer forms of

ethnography are incongruous with the needs of commercial designers – who typically demand rapid results from the fieldworker. Their approach clearly and deliberately goes against features such as intensity that define a prototypical ethnography. Their studies also exemplify some features of cognitive ethnography. For example, Blythin, Rouncefield and Hughes (1997) report an ethnographic study to inform changes in management practices within a large retail bank - providing a clear demonstration of purposiveness. In addition, these researchers are selective in their observational focus. For example, Hughes, O'Brien, Rodden & Rouncefield (1997) suggest that there are three key viewpoints that form the focus of ethnographic analysis for systems design: setting or ecology of work, workflow and social/organisational viewpoints. These viewpoints are specific in that they focus observations towards issues of distributed coordination, plans and procedures and awareness of work.

We find the concept of ethnography forwarded by Hughes and colleagues appealing in relation to applied research on human factors. We would further argue that recognising the value of impure ethnographic techniques is potentially liberating for researchers, since effort can switch away from attempts to accommodate the methodological edifice implied by the term ethnography toward a focus on developing the strengths of new techniques in field research. For example, Viller and Somerville (1997) have explored the potential of UML (Unified Modelling Language) as a notation for communicating the three viewpoints they focus upon in a form amenable to systems designers. This kind of development demonstrates how a pragmatic focus can advance the utility of an ethnographic approach.

The second high-level dimension on which ethnographies differ concerns the 'focus' of observation, with extremes placed in individual and social/organisational contexts. Again, much of the pioneering research on situated action and distributed cognition (e.g., Lave, 1988; Suchman, 1987; Hutchins, 1995) clearly fits well with the social/organisational focus. Indeed, this work presents strong theoretical arguments concerning the relation between cognition and culture. A key aspect of these arguments, exemplified by Hutchins (1995) in his studies of 'cognition in the wild', is the

establishment of an opposition between a traditional view of culture as simply affecting the internal cognition of individuals, and a view of cultural activity systems which have cognitive properties of their own that differ from those of the individuals who are members of such systems.

Research on cognitive activity in context has been invaluable in highlighting the real importance of the roles of culture, context and history in human cognition and the potential limitations that derive from maintaining a standard epistemology of distributed cognition as espoused, for example, by Larkin (1989) and Vera and Simon (1993). For example, the work of Blythin et al (1997) is aimed at developing multi-user systems to support the social and organisational aspects of banking practices. Thus their focus is upon social and organisational work issues, a notion captured neatly by the phrase 'sociality of work'. However, a shift towards a social/organisational focus needs to be balanced by retaining sight of the opposite end of the focus dimension: Few researchers would deny that individual cognitive phenomena exist, or that these are irrelevant to systems design. To maintain an explicit focus upon individual cognitive phenomena when adopting rich data collection approaches makes sense when the individual is the prime concern.

Furthermore, we suggest that the greatest value of the 'focus' dimension derives from comparing observations gleaned from each of the extremes. For example, Hutchins (1995) observes how the well known phenomenon of 'confirmation bias' in decision-making, a tendency to select verifying rather than falsifying observations (e.g., Baron, 1985), varies according to focus, with some team decision-making contexts showing an absence of the bias compared with some individual contexts. The example we describe later in this paper, of developing an indexing tool for supporting designers' ongoing work, required a focus upon the individual, while capitalising upon the benefits of an ethnographic approach to observation within a rich and realistic work environment. Like Hutchins, we provide an example of a well-known cognitive phenomenon, 'satisficing' (e.g., Simon, 1969), that is apparent in individual design work but appears to be 'managed out' of team contexts. What makes the dual focus of cognitive ethnography so

important to us is that the lessons learnt from its application allow the design of artifacts sensitive to both individual and social/organisational issues: We had to know about the individual design context to know what was missing from it, and we had to know about the team design context to know what to add to the tool for supporting designers working on their own. In the end, the 'surrogate manager' metaphor that we adopted for the indexing tool emerged directly from the comparative focus approach of cognitive ethnography.

Of the three features that define cognitive ethnography, none of the examples we have considered so far has addressed verifiability. In principle, there is no reason why any ethnographic data set cannot be subjected to some form of validation or verification, with varying degrees of objectivity. Indeed, an ethnomethodological stance places considerable emphasis on processes such as triangulation (e.g., Hammersley & Atkinson, 1993). Another approach adopted by ethnographers to verify their findings is grounded theorising (e.g. Strauss & Corbin, 1990), in which emergent properties of the data take on theoretical status and are then subjected to subsequent verification in further data analyses. Notwithstanding these efforts to validate observations, there is generally less emphasis upon verifiability in prototypical ethnography than in traditional cognitive methods used to study work activities. Indeed, features of a prototypical ethnography such as personalisation, reflexivity and self-reflection described above militate against objective verification. Our notion of cognitive ethnography places great weight upon verifiability, in part because, unlike grounded theory approaches, the feature of purposiveness places one within a hypothetico-deductive research tradition, of which objective empirical verification is the key. Furthermore, we see great benefits in adopting some of the specific methodologies used in cognitive research to triangulate observations made under an ethnographic method. This cross-method triangulation contrasts with, yet complements, the within-study data triangulation approaches advocated by ethnomethodologists. In the next section we overview a project that illustrates the goal and focus of a cognitive ethnographic approach, as well as showing how cross-method verification can take place.

Supporting design re-use: A prima facie case for cognitive ethnography

In this section we illustrate the use of cognitive ethnography by describing an investigation aimed at facilitating the re-use of design information. Re-use is the process by which previously encountered information such as solution options, problems, trade-offs, and evaluative critiques is used by designers in tackling a current problem. On the empirical side, there is evidence that previous design concepts and prototypes do get re-used extensively in routine design situations which are characterised by relatively well-defined design tasks where known solutions can be re-applied (e.g., Schön, 1988; Gero, 1990; Maiden & Sutcliffe, 1992; Visser, 1996). Our intuition is that there are at least two ways in which re-use could benefit designers in more innovative environments. First, the re-use or adaptation of existing design concepts, solutions and critiques makes cost-effective use of design efforts. Second, if designers know that ideas may be useful to future projects (even if not implementable immediately), then they may explore design options more fully and creatively. However, little is currently known about the nature, extent and efficacy of re-use in innovative design environments, though designers frequently tell us that re-use does not play an important role in innovative design practices (this is perhaps not surprising, in that the terms ‘re-use’ and ‘innovation’ are linguistically at odds each other).

We are currently engaged in a three-year project to specify and develop a computer-based system for indexing re-use information, called DESPERADO (DESIGN Process Encoding and Retrieval by Agent Designated Objects; Ormerod, Mariani, Ball & Lambell, in press). The project has two parallel strands. One is empirical, in which cognitive ethnography is being employed to investigate the re-use activities of design groups working in large-scale commercial design settings. The other strand is computational, in which concepts of object-oriented and intelligent agent systems are being applied in developing DESPERADO. Inherent in the project is the dependence of the computational work upon the outcomes of the empirical work, since our analysis of the literature suggests that previous attempts to foster re-use through technological

support have foundered upon the absence of a proper understanding of the contexts in which re-use activities occur.

Selecting the research method

Our first task was to determine an appropriate method for collecting data. In previous research, we have used verbal protocol analysis (Ormerod and Ball, 1993; Ball, Evans, Dennis & Ormerod, 1997) and longitudinal diary studies (Ball, Evans & Dennis, 1994) to investigate the cognitive processes underlying design problem-solving. Like other researchers (e.g., Guindon, 1990) we have focused upon individual designers, although the tasks undertaken by participants varied from single-session and relatively artificial tasks (Ormerod and Ball, 1993) to large-scale realistic tasks carried out over many months (Ball et al., 1994). These studies revealed a number of phenomena, some of which we discuss below, that served as the basis for initial speculations about the design of DESPERADO.

Whilst our research findings may appear sufficiently compelling to begin the specification of re-use tools (or at least they were to the sponsors of the DESPERADO project), it remains the case that we have only observed designers on an individual participant basis. Thus it is not necessarily the case that our observations will scale up to the operation of multi-participant design teams. The highly context-bound nature of problem-solving in design ultimately appears to derive from the sheer magnitude and complexity of modern technological design tasks (e.g., jet engine design, telecommunications design) together with the vast scale of the cost/benefit ratio associated with success or failure. What such factors mean in practice is that technological design activity typically takes place in well-resourced companies where multiple agents (both people and computer systems) work collaboratively on an artefact design task in a divide-and-conquer fashion. Within any company, design activity is typically guided by local organisational goals and managerial structures. Rejection of cognitive methods in favour of a form of ethnographic research method reflects the context-bound nature of the domain.

A pure form of ethnography was also deemed inappropriate, for the reasons outlined above. Notably, the principle aim of specifying an artefact made our data collection purposive from the outset. Also, the fact that DESPERADO had to meet the needs of many end-user companies meant that we had to capture high level generalities in re-use activities where apparent, as well as identifying company- or team-specific practices. Perhaps most importantly, we needed to be able to convince our sponsors, and more particularly the intended users, that we had strong theoretical grounds on which to develop the tool and an unambiguous way of evaluating its effects upon design team performance (or at very least, of ensuring that the tool does not do more harm than good). These considerations placed a strong emphasis upon the verifiability of our data.

The purposive nature of our data collection meant that it was necessary to determine, at least at in a general sense, the types of data that we were to collect in the studies. Specifying the data to be collected before beginning an observational study runs counter to a number of features of a prototypical ethnography, notably intensity, independence and to a certain extent, richness. Yet it was obvious to us that it was appropriate, if not essential, to do so in this project. For example, we had limited access, partly through company-imposed limitations (e.g., access to designers ‘off-task’; protection of copyright, and in one case the official secrets act, preventing access to all information sources). Also, we had to collect and analyse sufficient data within one year to complete the specification and prototyping of DESPERADO within three years. Thus, we had to be certain prior to commencement of the project that we could gain sufficient data to inform system design. Therefore, specificity of data collection was considered essential.

Determining the data to be collected

In proposing the construction of DESPERADO, we were driven initially by theoretical considerations deriving from previous cognitive research. For example, there is evidence for satisficing in the selection and evaluation of putative solution options. Individual designers frequently become fixated upon single solution ideas (usually derived from prior experience) rather than exploring alternatives in order to optimise choices (e.g.,

Ullman, Dietterich & Stauffer, 1988), even when they recognise that the solutions are less than satisfactory (Ball, Evans, & Dennis, 1994; Ball, Maskill & Ormerod, 1998). To address the problem of satisficing, Ball et al. (1994) argue that a system for promoting the retrieval of re-use information needs to facilitate the active consideration of alternative design solutions. A key aspect of this solution would be to provide a mechanism for system-initiated decisions (perhaps using intelligent agent technology; e.g., Shoham, 1993) as to what, when, and from where, to retrieve during an ongoing design process.

Data collection was also motivated by pragmatic considerations emerging from our own task analysis of the design process and also from previous experience of design re-use projects. A major problem in re-use relates to the encoding of design information. Handling the sheer mass and complexity of project information is often seen as prohibitively wasteful of designers' time (Conklin & Burgess-Yakemovic, 1991). Also, unless documentation about the limitations of previous designs is accessible, there is a risk that using old solutions can perpetuate sub-optimal designs. In overcoming such difficulties, DESPERADO should be able to assign labels to design information that define its nature and potential utility in a rich and multi-faceted manner. Thus, the development of a 'coding scheme' for labelling re-use information was a critical objective of the project.

A promising approach to defining design information classes, and one we adopted as a starting point for our coding scheme, is MacLean, Young, Bellotti and Moran's (1991) Questions, Options, and Criteria (QOC) formalism for design rationale. Questions describe key design issues, Options describe possible solutions to design problems, and Criteria describe approaches to assessing options. MacLean et al. found considerable overlap between the designers' discourse and the QOC concepts, though recent usability analyses of QOC have suggested extensions to the notational vocabulary (e.g., Buckingham Shum, 1996). It is noteworthy, however, that QOC research has concentrated on designers working within highly artificially-constrained design situations. Our task analysis of design re-use indicates that at least one level of notation

below QOC would be required to index design information in a re-usable way. However, determining a re-use coding scheme remains an empirical question that can only be answered through in-depth studies in fully dynamic and situated environments.

A second pragmatic consideration was the need to capture data from previous and ongoing projects in each of our end-user companies that could ‘seed’ DESPERADO. We regarded stocking the tool with re-use information appropriate for each of our end-user companies as vital, both to enable empirical evaluation of the tool, and also to secure its acceptance by its intended users, since we were aware that in imposing yet another technology upon designers we had to offer added value from the start, which would be missing in a tool that contained no re-useable data.

These theoretical and pragmatic considerations led us to use cognitive ethnography to collect data that address three key objectives:

- i. Identifying the data of re-use: Here we are seeking answers to the questions "What kinds of information are re-used?" and "What information that does not appear to be re-used currently might be amenable to re-use"? One specific outcome is to identify the most useful unit of re-use (subsequently referred to as an ‘episode’). A secondary outcome is to collect a set of episodes for seeding the re-use tool.
- ii. Exploring natural re-use: Here we are pursuing the questions "How do designers encode information for subsequent re-use?" and "How and when do designers retrieve information for re-use?" One outcome is to use evidence of existing re-use practices inform the development of a coding scheme. A further outcome is to attempt to extract regularities in natural re-use practices to inform the development of system-initiated retrieval.
- iii. Situating re-use in existing design practices: The question here is "What happens in the design process when re-use occurs?" For example, we need to investigate whether phenomena evident when we study individual designers, such as satisficing, are also shown in team design contexts.

As a consequence of these objectives, we adopted a deliberately constrained view of the data to be collected from the outset.

Applying the cognitive ethnographic method

In order to collect data that address these objectives, we carried out a short observational study at each of our four end-user companies, all major international leaders in their own fields of telecommunications, computer systems, aerospace and industrial design. Prior to each study, we held a series of meetings (three or four one-hour sessions over the course of a six-month period) with the managers and selected members (as determined by the end-user groups) of each design team, in order to set up the studies. This phase was basically a process of capturing background information to familiarise ourselves with the design domain of each end-user group as well as negotiating the terms under which, and when, our researcher could observe each group.

Our researchers spent between two and four weeks with each targeted design group, one researcher per group, and two researchers associated with the project. In each group, the researcher was a 'semi-participant' observer, in that he or she did not contribute to the design work of the teams directly but was frequently involved in design discussions, often acting as a sounding board, sometimes even offering advice, opinions and informal evaluations of design options. The researcher was given a desk in the main design room of each team (most teams were based in open-plan offices, with only minimal physical divisions between individual designer's desks), and was, generally speaking, allowed to observe the work of individual designers and sit in upon design team meetings whenever she wished, unless designers asked the researcher to withdraw or withheld documents from view (a relatively rare occurrence as it turned out). The researcher also conducted informal and unstructured interviews with design team members, usually on a one-to-one basis, though occasionally with two or more designers.

Data collection consisted of recordings (by audiotape or videotape where possible) of conversations, interviews and meetings, and field notes that detailed the activities observed by the researcher along with documentation (e.g., company standards, memos, design briefs, drawings and notes). These data were complemented by notes detailing the researcher's comments upon, and attitudes towards, the events and contexts

that were experienced¹. In total, the researchers collected around 150 hours of audio-and videotape data, along with 100 or so pages of field notes. Even with selective data gathering such as this, the data set is extremely large (a year later we are still completing our analyses).

Some highlights of the resulting data

Partly because of the sheer quantity of data and the limited length of this article, the review of research findings presented here is necessarily selective. We focus upon how the outcomes of the cognitive ethnography informed our three key objectives outlined above.

i. Identifying the data of re-use: We began our analysis by using the QOC scheme (MacLean et al., 1991) to guide the sectioning of data into re-use ‘episodes’. In this sense, our analysis was driven by the goal of identifying the nature and size of the ‘chunk’ of a re-use information appropriate for handling within DESPERADO.

It turned out from inspection of the transcripts and discussion with the researcher that shifts in question focus were an apparently natural transition point between episodes. The size of episode did vary both within and between companies. For example, data from the telecommunications and industrial design groups indicate much longer episodes than data from the aerospace group. This may in part reflect the different types of design work focused upon in each company (e.g., review-meetings and focused team design work versus individual or incidental collaborative), but it is also representative of the kinds of design activity that the groups were typically involved in (notably the scale of questions pursued by each design group). The sectioning of episode by question also, perhaps fortuitously, gave us a manageable chunk size for indexing re-use information. For example, two weeks of data collection from the aerospace company gave us 50 episodes that have subsequently been used to seed the DESPERADO database for that company.

¹For example, the field notes taken from one company made frequent reference to the 'Nomad', an evocative name given by the researcher to an individual who seemed to spend his time wandering around the design group without any specific role or agenda, and whose influence upon design activities was not always welcomed by other members of the design team.

In order to investigate the validity of the ‘episode by question’ notion further, the second author and one of the researchers carried out independent sectioning by question of episodes taken from the first week of the study at the aerospace group. The grain size of episodes was surprisingly similar (TO’s mean episode length = 151 words; The researcher’s mean episode length = 150 words). The variance of the two sets of values was also very similar. There was remarkable correspondence between the episodes that both coders identified. However, TO’s sectioning revealed twice as many episodes (34) compared with the researcher’s (17), and 18 of TO’s episodes were not identified as episodes by the researcher. Close scrutiny of these 18 episodes suggests that they do not really convey design work as such, but are instances of designers presenting background information to the researcher in order to furnish her with an understanding of the project’s aims and context etc. This provides an interesting illustration of the value of an ethnographic approach to the analysis of the data: in essence, a participant observer has a closer relationship to the data than an independent observer, and is better placed to interpret the significance of emerging observations. As a consequence, the final sectioning of data into episodes was based mainly upon episodes that were recognised by both a researcher and one of the authors. There were very few instances in the data where the researcher recognised an episode that was not identified by one of the authors. Once episodes were identified, they were coded in order to identify the focus of design work adopted by the individual designer. The scheme consisted of a set of codes that identify functional, structural, procedural and optimisation-oriented views. The coding of episodes under this scheme is illustrated in Figure 1.

 Insert Figure 1 about here

We then met to review the episodes in order to develop a second coding scheme for capturing the contextual nature of each episode. Figure 2 shows an example of an episode coded under this context-oriented scheme. The role of the researchers was

central to the development of the scheme, since they were able to provide insights into the nature of design practices that were not available to the authors. For example, whilst the authors noted that episodes tended to have particular players (individuals or groups of designers) associated with them, the researchers were able to point out important distinctions in the roles of different players. In particular, some participants are identified as ‘owners’ of the particular project that an episode concerns. In the aerospace group, ‘ownership’ is an explicit managerial practice intended to foster the maintenance of project information and skill over time. In essence, the project owner is the principal source of information about that project within the company, and has responsibility for maintaining knowledge about the project both during and after project completion.

In the industrial design group, there appeared to be tacit use of the ownership notion, in that each project had an individual who was recognised by other designers as the ‘project champion’, who carried with him or her the knowledge base necessary to guide the project, and who appeared to the researchers to be ‘carrying the torch’ for a particular project. We are currently in the process of identifying and validating further player roles that appear to be important (e.g., there is evidence in one company for the adoption of a ‘devil’s advocate’ by one or more players in any collaborative setting, whose self-selected role seems to be to critique the work of others: we are examining transcripts from other companies to establish whether this is a general pattern in design groups or whether it is particular to this company).

 Insert Figure 2 about here

Our analysis to date has provided us with a coding scheme that has two main types of category. The first is a set of categories that describe the data of each episode (the players, the project names and types, the components, artefacts and tools that are discussed, and the antecedent and consequent episodes associated with the current episode). These categories are generally speaking explicit in the episode transcripts. The second set of categories is orthogonal descriptions that capture aspects of design activity.

These include the stage (e.g., requirements specification, conceptual design, detailed design, etc.), process (e.g., the generation, evaluation and simulation of putative solution options or critiques) and scope (e.g., project-specific, general to design practices, organisational, managerial, company policy-related, etc.) of each episode.

Our application of the coding scheme also revealed that each episode contained a constellation of options and criteria associated with each question. For example, in the aerospace data, 64% of the episodes contain at least one option and criterion associated with the episode question, there were two episodes in which a criterion was examined in the absence of any design options, but there were no episodes in which options were discussed in the absence of criteria.

ii. Exploring natural re-use: Each of the four companies had in place a number of existing central information repositories, such as paper-based archives, computer-based CAD (computer-aided design) databases and diary-based file stores. Despite the existence of these repositories, in addressing the question "How do designers encode information for subsequent re-use?" we were surprised to see little evidence of systematic encoding of design information. Storage was typically left to the 'owners' of each project, who tended to place only a sample of official documents (e.g., requirements specifications, tenders, CAD drawings) in shared repositories, the majority of documentation being held in individual designer's computers and filing cabinets. Information was typically encoded (i.e. labelled) in a piecemeal fashion, with file names generally specifying either the players, project names, components under design, the date, or some combination of these categories. There was no evidence of encoding by process (i.e., of a sort captured by our orthogonal categories) or by focus (i.e., the QOC categories).

The absence of systematic encoding practices was, in our evaluation, reflected in the ways, and the extent to which, information was retrieved for re-use. The extent to which information repositories in the companies were used successfully was extremely variable, though the overall evaluation of the systems (not only by ourselves but more importantly by designers and design group managers who used the systems) emerging

from our studies is that they failed to support systematic and optimal information re-use. For example, one company had installed a project management system based around a computer-networked diary, which automatically created folders in a project file hierarchy for designers to deposit project-related information. In demonstrating the system to our researcher, the team manager was unable to find a single folder that had any contents, despite the fact that the system had been running for two years prior to our visit. It became clear, from subsequent conversations with members of the design team, that they maintained their own information repositories, only occasionally shifting files to the designated project folders immediately prior to team meetings. In another example, the researcher observed one designer engaging in a week-long search for a specific piece of design information. It took him a day to locate the 'owner' of the project to which the information pertained, and a further two days to find that this individual had in fact retired from the company some years ago, taking with him the knowledge of how to retrieve the information sought by the designer (though he was sure that it was held "somewhere in the central files of the group"). These examples are extreme, though they illustrate a tendency among the designers to place little reliance on retrieval from central repositories.

Despite the limited evidence for successful encoding and retrieval using the central repositories, there was strong evidence for the frequent re-use (or attempted re-use) of design information from past projects to inform ongoing work. Two main classes of re-use activity were observed: the first, and most common one, was for individual designers to recall from memory previous work in which they were involved or components and solution options (and very occasionally, critiques) that they had experience of. The second class of re-use was essentially managerial, in that design team managers were often observed making suggestions, either to individuals or during team meetings, as to what kinds of information designers might seek to solve particular problems as well as suggested ways of locating relevant information.

iii. Situating re-use in existing design practices: A key issue motivating our cognitive ethnography was to understand the extent and consequences of sub-optimal re-

use strategies such as satisficing. The results are surprising in the degree to which they contrast so markedly with our previous studies of individual and de-contextualised designers (e.g., Ball et al., 1998). Almost all episodes appear to reflect a motivated attempt by designers to generate and evaluate multiple solutions options (albeit ones typically stemming from their own memories) with the aim of maximising the choice of a best alternative in relation to a particular design question. The evaluation of design options frequently takes the form of sophisticated trade-off analyses, with options being systematically pitted against each other in terms of key design criteria. Importantly, too, this evidence for optimising attempts at multiple solution search and evaluation is a finding that generalises across the different companies we researched and the diverse design tasks that were being tackled within these companies.

The prevalence of highly effective search and evaluation strategies in these companies, coupled with evidence for less effective strategies in de-contextualised individuals, suggests that re-use behaviours may be closely determined by the aspects of the context in which design work is occurring. For example, a major function of design review meetings was for the design team to explore and critically appraise a range of alternative design concepts in an in-depth manner. Indeed, the team dynamics are highly tuned toward explicating novel design options, new evaluation criteria, unconsidered design constraints and the like. Quite significantly, too, some more experienced project champions appear to enter the design review meeting well armed with a range of alternative design options for discussion together with detailed knowledge of their associated costs and benefits. This kind of pre-emptive approach to design review appears to be a sophisticated strategy aimed at dealing effectively with the penetrating questions that arise from those adopting the devil's advocate role within the team.

A final and striking feature of the team dynamics is the role played by the team manager as a safeguard against premature commitment to single solution options should either the project champion or the group of collaborating designers become overly fixated upon a particular idea or concept prior to a full and cautious exploration of viable alternatives. It appears that the team manager has a very clear and explicit notion of the

importance of ensuring multiple solution search and evaluation such that they are seen to perform an important steering function when designers get fixated upon single or unsatisfactory solution ideas or need support in pursuing alternatives. So, for example, in one team design session, where a project champion is describing the unsatisfactory aspects of a solution that they are actually quite committed to, the team manager interjects to say “I think the thing to do is look at all the other options”. This simple interjection is particularly striking as it is only the third statement that the manager has made in the first 30 minutes or so of the team design session and it therefore appears to carry considerable weight. Another example, from a different company, again relates to a situation where a member of the design team is beginning to become overly fixated upon the negative aspects of one solution option that has been suggested. The project manager strives to re-inforce the need to explore fully the available options by saying ‘I think we would be mistaken if we assumed it was one or other answer . . . I think we agreed in the last meeting to look at the two extremes’. This seems to be sound advice given that the team is in the very early, conceptual stages of the design process where options need to be fully explored.

Verifying the outcomes of the ethnography

The preceding description demonstrates the purposive and specific nature of data collection in cognitive ethnography. The third aspect, verifiability, is a critical part of cognitive ethnography, since it is only possible to gain the confidence of clients and investors by being able to demonstrate that ones 'story' has some degree of objectivity, ubiquity and salience. In our study, we undertook three approaches to verifying our observations: first, the application of a coding scheme to the data for cross-company/session comparisons (as illustrated in Figure 2); second, the embodiment and evaluation of observations in a computer-based system (as described in Ormerod et al., in press), and third, comparisons of empirical outcomes across cognitive and ethnographic methods. We focus here upon the third, since the mixing of research methods across what are generally seen as radically different, even antagonistic,

research traditions provides what we believe is one of the most useful features of cognitive ethnography.

The combination of ethnographic approaches with traditional controlled methods of cognitive psychology for verifying findings is illustrated by a study carried out with the team of software designers described earlier (Ormerod, Rummer & Ball, in press). They took the coded data from the ethnographic study conducted in this company, and isolated sets of design episodes. These were subsequently used as materials for a study using a sort methodology (e.g., Hoffman, Shadbolt, Burton & Klein, 1995). The designers categorised the episodes into natural categories that grouped them in ways that were personally meaningful. The aims of the study were, first, to inform the development of the indexing system by eliciting a set of episode categories, and second, to triangulate observations made in the ethnographic data, both at this company and at other sites. Analysis of the data (using cluster analysis and multiple regression) revealed three key results. First, there was a remarkable degree of consistency amongst participants in category generation and assignment. Second, this consistency was achieved largely because designers used an impoverished set of dimensions for categorisation (classifying episodes by artifact, stage, design task and marketing issue). Third, the order in which categories are generated is related to the role of the designer in the team (e.g., manager, product champion, and implementor). The study is novel, in that it is the first we are aware of in which participants categorise materials they themselves have produced. It also provides an objective verification for many of our observations drawn from the ethnographic data.

Conclusions

In this paper we have argued for a limited, yet definable, use of ethnographic methods for human factors research, which is distinguished from pure ethnography in terms of

purposiveness, verifiability and observational specificity. On the basis of the analyses we conducted, we have drawn up a requirements specification and have programmed a prototype interface for the DESPERADO indexing tool. In particular, our use of cognitive ethnography has informed the development of an episode model that provides a metaphor for the systematic encoding of design information, and has enabled the elicitation and validation of a coding scheme for labelling information for subsequent re-use that captures both natural re-use based upon recall from memory of project data-related aspects, and also the prescriptively optimal yet less apparently natural retrieval through application of process-oriented retrieval rules. Our next step is take the prototype tool to the design groups and allow them to test and critique the tool. We are confident that we have got much further in systems design than a full-blown ethnographic study would have allowed in the same timescale.

There are a number of issues which space precludes us from addressing in which the notion of cognitive ethnography might be critiqued. For example, although our researchers were to some extent participant observers, their role was nonetheless relatively neutral with respect to the design teams that they observed. This can be contrasted with the view taken by some proponents that a pure ethnography requires the integration of the researcher into the group under study. It was clearly not possible for our researchers to become integrated into the design teams, since all the designers were highly qualified and experienced specialists in their own fields, and it was clearly not feasible for us to train our researchers to a level of competence in each design domain. Indeed, even acquiring sufficient knowledge to understand the basic aspects of each design domain was in itself a challenging task which took at least three weeks for each company studied.

The fact that we chose to study four companies for short periods of time rather than, say, one company for a more extensive period might be questioned by some ethnographers. However, as well as providing data that inform the development of more general support tools than would be allowed by studying only a single company's practices, the cross-company comparisons that our data allowed also gave us a way of

verifying the outcomes from each study. Ethnographers might argue that their observations achieve some validation by their recurrence, and that anyway some of the most significant and informative outcomes of an ethnographic analysis are the non-recurring and distinctive behaviours that are observed. However, it is clear to us that both our sponsors, and perhaps more importantly the companies themselves, would be more convinced by findings that can be verified in to some degree of objectivity. This is particularly important given that we are proposing to impose upon designers another information handling tool despite the evidence that they see of the failures of existing information repositories.

Whilst our notion of cognitive ethnography is cognitivist in its orientation, the use of the technique also affords the advantages offered by adopting an ethnographic stance. Indeed, cognitive psychologists and systems designers should be confident in adding ethnography to their repertoire of techniques. For example, an important finding that demonstrates the advantages of the use of the method over more traditional cognitive methods for human factors research is the evidence for systematic design practices that avoid problems shown by individual designers working in de-contextualised studies. The study of complex human activity in general, and human factors in particular, is littered with arguments, and in our opinion false dichotomies, concerning the validity of different research methodologies. We offer cognitive ethnography as an initial attempt at a rapprochement between ethnographic and cognitive methods, in the hope that a demonstration of combining the strengths from each approach can be made.

References

- BALL, L.J., EVANS, J.St.B.T., & DENNIS, I. (1994). Cognitive processes in engineering design: A longitudinal study. *Ergonomics*, **37**, 1753-1786.
- BALL, L.J., EVANS, J.St.B.T., DENNIS, I., & ORMEROD, T.C. (1997). Problem-solving strategies and expertise in engineering design. *Thinking and Reasoning*, **3**, 247-270.
- BALL, L.J., & ORMEROD, T.C. (1995). Structured and opportunistic processing in design: A critical discussion. *International Journal of Human-Computer Studies*, **43**, 131-151.
- BALL, L.J., MASKILL, L., & ORMEROD, T.C. (1998). Satisficing in engineering design: Causes, consequences and implications for design support. *Journal of Automation in Construction*, **7**, 213-227.
- BARON, J. (1985). *Rationality and intelligence*. Cambridge: Cambridge University Press.
- BASZANGER, I., & DODIER, N. (1997). Ethnography: Relating the part to the whole. In D. SILVERMAN (Ed.), *Qualitative research: Theory, method and practice*. (pp. 8-23). London: Sage.
- BLYTHIN, S., ROUNCEFIELD, M., & HUGHES, J.A. (1997). 'Never mind the ethno' stuff – What does all this mean and what do we do now?': Ethnography in the commercial world. *Interactions*.
- BUCCIARELLI, L.L. (1988). An ethnographic perspective on engineering design. *Design Studies*, **9**, 159-168.
- BUCKINGHAM SHUM, S.J. (1996). Analysing the usability of a design rationale notation. In T.P. MORAN and J.C. CARROLL (Eds.), *Design rationale: Concepts, techniques, and use*. (pp. 185-215). Hillsdale, NJ: LEA.
- CARROLL, J.M., & CAMPBELL, R.L. (1989). Artefacts as psychological theories: The case of human-computer interaction. *Behaviour and Information Technology*, **8**, 247-256.

- CONKLIN, E.J., & BURGESS-YAKEMOVIC, K.C. (1991). A process-oriented approach to design rationale. *Human-Computer Interaction*, **6**, 357-391.
- D'ANDRADE, R.G. (1990). Some propositions about the relations between culture and human cognition. In J.W. STIGLER, R.A. SHWEDER, and G. HERDT (Eds.), *Cultural psychology*. Cambridge: Cambridge University Press.
- ERICSSON, K.A., & SIMON, H.A. (1993). *Protocol analysis: Verbal reports as data (Revised Edn)*. Cambridge, MA: MIT Press.
- GERO, J.S. (1990). Design prototypes: A knowledge representation schema for design. *AI Magazine*, **11**, 26-36.
- GOODWIN, C., & GOODWIN, M.H. (1995). Formulating planes: Seeing as situated activity. In D. MIDDLETON, and Y. ENGESTROM (Eds.), *Cognition and communication at work*. Cambridge: Cambridge University Press.
- GUINDON, R. (1990). Designing the design process: Exploiting opportunistic thoughts. *Human-Computer Interaction*, **5**, 305-344.
- HAMMERSLEY, M., & ATKINSON, P. (1983). *Ethnography: Principles in practice*. London: Routledge.
- HOFFMAN, R. R., SHADBOLT, N. R., BURTON, A. M., & KLEIN, G. (1995). Eliciting knowledge from experts: A methodological analysis. *Organizational behavior and human decision processes*, **62**, 129-158.
- HOFFMAN, R.R., FELTOVICH, P.J., & FORD, K.M. (1997). A general framework for conceiving of expertise and expert systems in context. In P.J. FELTOVICH, K.M. FORD, and R.R. HOFFMAN (Eds.), *Expertise in context*. (pp. 453-580). Cambridge, MA: MIT Press.
- HUGHES, J.A., KING, V., RODDEN, T., & ANDERSEN, H. (1994). Moving out from the control room: Ethnography in system design. In *Proceedings of CSCW '94*, Chapel Hill, North Carolina.
- HUGHES, J.A., O'BRIEN, J., RODDEN, T., & ROUNCEFIELD, M. (1997). Ethnography, communication and support for design. In C. HEATH, and P. LUFF (Eds.), *Workplace studies*.

- HUTCHINS, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- LAKOFF, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- LARKIN, J. (1989). Display-based problem solving. In D. KLAHR, and K. KOTOVSKY (Eds.), *Complex information processing: The impact of Herbert A. Simon*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- LAVE, J. (1988). *Cognition in practice*. Cambridge: Cambridge University Press.
- LAVE, J., MURTAUGH, M., & DE LA ROCHA, O. (1984). The dialectic of arithmetic in grocery shopping. In B. ROGOFF, and J. LAVE (Eds.), *Everyday cognition: Its development in social context*. Harvard: Harvard University Press.
- LUNIN, L.F., & FIDEL, R. (1994). Perspectives on indexing. *Journal of the American Society for Information Science*, **45**, 570-627.
- MACLEAN, A., YOUNG, R.M., BELLOTTI, V.M.E., & MORAN, T.P. (1991). Questions, Options, and Criteria: Elements of design space analysis. *Human-Computer Interaction*, **6**, 201-250.
- MAIDEN, N., & SUTCLIFFE, A. (1992). Analogously based reusability. *Behaviour and Information Technology*, **11**, 79-98.
- NISBETT, R.E., & WILSON, T.D. (1977). Telling more than we can know: verbal reports on mental processes. *Psychological Review*, **84**, 231-259.
- ORMEROD, T.C., & BALL, L.J. (1993). Does programming knowledge or design strategy determine shifts of focus in Prolog programming? In C.R. COOK, J.C. SCHOLTZ, and J.C. SPOHRER (Eds.), *Empirical studies of programmers: Fifth workshop* (pp. 162-186). Norwood, NJ: Ablex.
- ORMEROD, T.C., MARIANI, J.A., BALL, L.J., & LAMBELL, N.J. (in press). Desperado: Three-in-one indexing for innovative design. To appear IN M.A. SASSE AND C. JOHNSON (Eds.), *Proceedings of the Seventh IFIP Conference on Human-Computer Interaction—INTERACT '99*. London: IOS Press.
- ORMEROD, T.C., RUMMER, R., & BALL, L.J. (in press). An ecologically-valid study of categorisation by design experts. To appear in D. HARRIS (Ed.), *Engineering*

- Psychology and cognitive ergonomics—Volume four*. Hampshire: Ashgate Publishing Ltd.
- PAHL, G., & BEITZ, W. (1984). *Engineering design*. London: The Design Council. [Original German text, 1977, Berlin: Springer-Verlag].
- RACHEL, J. (1996). Ethnography: Practical implementation. In J.T.E. RICHARDSON (Ed.), *Handbook of qualitative research methods for psychology and the social sciences*. (pp. 113-124). Leicester: British Psychological Society.
- ROUNCEFIELD, M., HUGHES, J.A., & O'BRIEN, J. (1997). Ethnography: Some practicalities of ethnographic analysis. Technical Report (Ref: CSEG/27/1997), Computing Department, Lancaster University, UK.
- SCHÖN, D.A. (1988): Designing: Rules, types and worlds. *Design Studies*, **9**, 181-190.
- SHOHAM, Y. (1993). Agent-oriented programming. *Artificial Intelligence*, **60**, 51-92.
- SIMON, H.A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- STRAUSS, A., & CORBIN, J. (1990). *Basics of qualitative research: Grounded theory procedures and techniques*. London: Sage.
- SUCHMAN, L. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge: Cambridge University Press.
- TAYLOR, M. (1994). Ethnography. IN P. BANISTER, E. BURMAN, I. PARKER, M. TAYLOR, and C. TINDALL (Eds.), *Qualitative methods in psychology: A research guide*. (pp. 34-48). Buckingham: Open University Press.
- TOREN, C. (1996). Ethnography: Theoretical background. In J.T.E. RICHARDSON (Ed.), *Handbook of qualitative research methods for psychology and the social sciences*. (pp. 102-112). Leicester: British Psychological Society.
- ULLMAN, D.G., DIETTERICH, T.G. & STAUFFER, L.A. (1988). A model of the mechanical design process based on empirical data. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, **2**, 33-52.
- VERA, J., & SIMON, H.A. (1993). Situated action: A symbolic interpretation. *Cognitive Science*, **17**, 7-48.

- VILLER, S., & SOMMERVILLE, I. (1997). Coherence: An approach to representing ethnographic analyses in systems design. Technical Report (Ref: CSEG/7/1997), Computing Department, Lancaster University, UK. [To appear in this IJHCS special issue on 'understanding work and designing artefacts'.]
- VISSER, W. (1990). More or less following a plan during design: Opportunistic deviations in specification. *International Journal of Man-Machine Studies*, **33**, 247-278.
- VISSER, W. (1996). Use of episodic knowledge and information in design problem solving. In N. CROSS, H. CHRISTIAANS, and K. DORST (Eds.), *Analysing design activity*. (pp. 271-289). Chichester: Wiley.
- WOOLGAR, S. (1991). Configuring the user: The case of usability trials. In J. LAW (Ed.), *A sociology of monsters: Essays on power, technology and domination*. (pp. 57-99). London: Routledge.

Author Note

This research was supported by an ESRC Cognitive Engineering Programme grant (No: L127251027). We are grateful to our grant co-holder (John Mariani) and to Louise Maskill and Gary Spiers for their contributions to aspects of the reported work. We also thank the designers and managers in our end-user companies for their co-operation, and Bob Fields, Peter Wright, Robert Hoffman and two anonymous referees for comments on an earlier draft of this paper.

Correspondence concerning this article should be addressed to the first author. Electronic mail may be sent via L.J.Ball@derby.ac.uk

Figure legends

Figure 1. Example of episodes inferred from ongoing design work, showing two of eleven sub-functions pursued by the designer, in which he deals with the handling of keys by an automated car rental facility. Line numbers indicate the design level of each episode and their hierarchical organisation. Q, O and C refer to questions, options and criteria. Labels F (functional), S (structural), and P (procedural) describe the view the designer adopts with each episode element. Numbers in parentheses refer to the order of episode elements in the transcript, and ellipses refer to episodes collapsed for brevity.

Figure 2. Example of an episode coded according to its context.

1. GENERAL FUNCTIONAL CONCEPT

Q (F) What is the machine going to have to do? [8]

O (F) It has to print receipts, take credit cards, dispense and receive keys. [9]

1.3 DISPENSING KEYS

1.3.1 Q (P) For car rental when should keys be dispensed? [12]

O (P) At the end of interaction sequence. [13]

1.3.2 Q (F) How should keys be dispensed? [18] [36] [50]

O (S) Row of keys for each car-type ".. it would be similar to the way that chocolate machines work, it drops off next key for that kind of car.." [37] "If you use the chocolate machine idea, you've got a spiral that turns and the keys drop off into the hatch.." [50]

1.3.3 Q (F) Identification of dispensed keys... [38-40]

1.4 RETURN OF KEYS

Q (F) How does the machine deal with return of keys? [28]

O (F) Deal with key return as a separate function. [29]

1.4.1 Q (F) How to validate key return... [34-35]

1.4.2 Q (F) How does machine sort dispensed keys when returned? [41]

O1 (F) Keys manual-sorted daily and put back into dispenser. [47]

O2 (S) Standard keys fit all cars, have immobiliser with unique code. Specific code on the receipt. All keys are same on return. [48]

C Segregation of keys. [49]

1.4.3 Q (F) How does machine collect dispensed keys on return. [42]

O1 (S) Collect keys through hatch, they drop out of reach. [43]

O2 (S) Keys get put back on the other end of dispenser spiral. [51]

1.4.3.1 Q (P) When should the hatch be activated?

O1 (P) Hatch activated after removing credit card. [44]

O2 (P) Hatch activated before card return. [45]

C Prevent taking away of keys. [46]

Coding

Category	Sub-Category	Data
Transcript	Identifier with line refs	e5
Time		1.15pm-2.45pm, 13/1/97
Players		Ethnographer, Team leader (AS), Designer (HG), three other people from Compression Group
Company		Aerospace plc.
Episode type		Design Meeting
Scope	Project	266 engine
Focus	Question-oriented	How to alter performance requirements
	Option-oriented	1. Development of the new engine 2. Scaling an old engine
	Criterion-oriented	1. Performance capabilities 2. Forecasted changes in market requirements 3. Limited resources within the compressor group 4. Different applications -> differences in stress levels within the engine
Process	Clarification	√
Stage	Requirements Specification	√

Ethnographer's notes

The meeting moves on to a discussion of the 266 project. An economic study of the new 266 engine concept has been undertaken, and has come up against problems involving the development of the new engine versus scaling an old engine. The study also worked out figures to do with performance capabilities and forecasted changes in market requirements. The new figures from the study mean that some tweaking of the performance requirements is needed, and some compressor work needs to be done with the new parameters. There are limited resources within the compressor group, which means that they will have to concentrate their efforts on the military application rather than the civil. The different applications of technology also means that there are differences in stress levels within the engine. The normal civil application has a lower stress level than the civil/military dual application. However, the new combined engine

has technology standard to a previously defined engine. Many problems can be dealt with by making changes to the existing specification.