

Applying ethnography in the analysis and support of expertise in engineering design

Linden J. Ball, Cognitive and Behavioural Sciences Research Group,
Division of Psychology, University of Derby, Mickleover, Derby, DE3
5GX, UK

Thomas C. Ormerod, Department of Psychology, Fylde College,
Lancaster University, Lancaster, LA1 4YF, UK

Whilst many contemporary studies of design have claimed to be using ethnographic methods, the techniques which have been employed often diverge from the characteristics of pure ethnography as used in traditional anthropological and sociological research. We argue that this is entirely appropriate for applied ethnography that is conducted in pursuit of explicit design goals. In this paper, we explore the relationship between pure and applied ethnography, and their use in cognitive and social research. We also discuss how the outcomes of applied ethnography can be applied to i) the design of a computer-based design support tool and ii) the development of controlled experimental studies that retain sufficient ecological validity to capture realistic design expertise. We additionally argue that objectivity in empirical studies of design can be obtained only by triangulating observations across methodologies that embrace both ethnographic and laboratory-based traditions. © 2000 Elsevier Science Ltd. All rights reserved

Keywords: ethnography, engineering design, design tools, design rationale, research methods

1 Anderson, R J 'Representations and requirements: The value of ethnography in system design' *Human-Computer Interaction* Vol 9 (1994) pp 151–182

2 Forsythe, D E 'Using ethnography in the design of an explanation system' *Expert Systems with Applications* Vol 8 (1995) pp 403–417

3 Sommerville, I, Bentley, R, Rodden, T and Sawyer, P 'Cooperative systems-design'

The use of ethnographic methods in design research has proliferated in recent years^{1–4} since the seminal research of Suchman⁵ and Lave⁶ which affirmed the importance of studying real-world problem-solving and planning behaviours as situated practices. It often appears, however, that there is only a limited overlap between the approaches taken by design researchers employing what they term ethnographic methods and the ethnographic techniques that have formed the basis of much traditional anthropological and sociological research. In this paper we argue that this



www.elsevier.com/locate/destud

0142-694X/00 \$ - see front matter *Design Studies* 21 (2000) 403–421

PII: S0142-694X(00)00009-0

© 2000 Elsevier Science Ltd All rights reserved Printed in Great Britain

disparity is an inevitable consequence of the *purposes* for which ethnographic methods are being used. For example, the pursuit of a relatively pure form of ethnography would seem entirely appropriate in anthropological or sociological studies in which the focus is that of gaining insights into previously unexplored cultures or undocumented social practices. In design contexts, however, it is apparent that the goals of research tend to be more applied in nature, such as attempting to understand design behaviours in order to make design productivity more effective (e.g., through computer-based support or changes to existing organisational practices). It is these variations in purpose which we believe motivate ethnographic studies of design to differ—at times quite markedly—from more traditional ethnographies.

In the remainder of this paper, we first briefly review the characteristics of pure ethnography as espoused by those using the method in anthropological and sociological research, and discuss how applied ethnography borrows from some of these characteristics whilst deliberately challenging others. The arguments that we make in this section build upon points we have made⁷ where we introduce the possibility of developing a ‘cognitive ethnography’ of design. In subsequent sections of the paper we describe how we have employed ethnographic methods in two contexts: first in studying the processes of design re-use in a major aeronautics company in order to inform the development of a design-support tool, and second, in collecting materials for subsequent cognitive experimental research into design expertise in a major software development company. The focus of this paper is methodological rather than empirical, so we do not offer lengthy descriptions of ethnographic data and their interpretation, instead choosing to focus upon the potential strengths of a clearly specified applied ethnographic method.

1 ‘Pure’ ethnography versus ‘applied’ ethnography

There is no doubt that ethnography has the potential to play a valuable role in design research. First, the complex, ongoing and multi-faceted nature of commercial design projects excludes the use of experimental research methods since such factors do not readily lend themselves to experimental control. Second, the fact that much real-world design is situated within a rich and dynamic social context not only adds new complexities to any attempt to achieve experimental rigour, but also obviates the use of non-experimental techniques such as verbal protocol analysis⁸ which are solely geared toward understanding *individual* behaviours in socially-impoverished contexts. Third, the observation⁹ that much important decision making and negotiation within organisations takes place through incidental meetings outside of the nominal work area (e.g., in canteens, corridors and

Computer Journal Vol 37 (1994)
pp 357–366

4 Lloyd, P and Deasley, P *Ethnographic description of design networks* Paper presented at the First International Symposium on Descriptive Models of Design, Istanbul Technical University, Taksim, Taskisla, Turkey (1996)

5 Suchman, L *Plans and situated actions: The problem of human-machine communication* Cambridge University Press, Cambridge (1987)

6 Lave, J *Cognition in practice* Cambridge University Press, Cambridge (1988)

7 Ball, L J and Ormerod, T C ‘Putting ethnography to work: The case for a cognitive ethnography of design’ *International Journal of Human-Computer Studies* (in press)

8 Ericsson, K A and Simon, H A *Protocol Analysis* (2nd Edn) MIT press, Cambridge MA (1993)

9 Boden, D *The business of talk: Organizations in action* Polity Press, Cambridge (1994)

doorways) means that any empirical technique should have some capacity to track such *ad hoc* encounters as well as the more standard, scheduled interactions.

The importance of ethnography as an emerging method in design research is clearly attested to by the burgeoning number of published studies purporting to have used ethnographic methods. We have, however, argued elsewhere⁷ that such studies are often using a methodology that is distinct from the kinds of ethnography employed in traditional anthropological and sociological research, whilst still retaining some of the key features of such ethnographies. In support of this view, we have proposed that ethnography may best be understood as a *radial category* which exhibits *prototype effects*¹⁰. Radial categories encompass a central, prototypical case as well as a number of variations from this prototype. The prototypical case can then be viewed as a template against which one may examine other instances in order to assess their goodness-of-fit to the prototype. Under this conceptualisation it is possible for a case of ethnography to differ from the prototype in terms of qualitative or quantitative omissions in relation to one or more attributes of the prototype.

Our starting point for defining a prototypical case of ethnography has involved systematically examining the relatively pure forms of ethnography as espoused in standard textbooks and anthropological monographs^{11,12} as well as research papers discussing this method^{13–16}. From this literature review we have found that there appear to be ten consistently occurring features which exemplify a prototypical ethnography (see Table 1). Note that the category is radial in the sense that we would not expect every instance of an ethnography to conform precisely to these listed features but simply to retain the flavour of a majority of them. Such prototypical ethnographic methods are very closely identified with those of traditional anthropologists whose goal is primarily that of developing a personalised understanding of the rituals and customs of a society, culture or organisation for the advancement of knowledge *per se*. In such pure ethnographies the levels of subjectivity, reflexivity and self-reflection are high, observational openness and independence are paramount, and the intensity is extreme (some anthropological ethnographies last many years).

What is noticeable, however, when one examines existing design research which has purported to use ethnographic techniques, is that almost all such studies violate a significant subset of the characteristics of a prototypical ethnography, either by omission or by satisfying certain characteristics to only a weak degree. Two examples of published research which claim to be using ethnographic techniques may serve to illustrate such contra-

10 Lakoff, G *Women, fire, and dangerous things: What categories reveal about the mind* University of Chicago Press, Chicago (1987)

11 Hammersley, M and Atkinson, P *Ethnography: Principles in practice* Routledge, London (1983)

12 Bernard, H R *Research methods in anthropology: Qualitative and quantitative approaches* (2nd Edn) Sage, London (1995)

13 Taylor, M 'Ethnography' in **P Banister, E Burman, I Parker, M Taylor and C Tindall** (eds) *Qualitative methods in psychology: A research guide* Open University Press, Buckingham (1994) pp 34–48

14 Toren, C 'Ethnography: Theoretical background' in **J T E Richardson** (ed) *Handbook of qualitative research methods for psychology and the social sciences* British Psychological Society, Leicester (1996) pp 102–112

15 Rachel, J 'Ethnography: Practical implementation' in **J T E Richardson** (ed) *Handbook of qualitative research methods for psychology and the social sciences* British Psychological Society, Leicester (1996) pp 113–124

16 Baszanger, I and Dodier, N 'Ethnography: Relating the part to the whole' in **D Silverman** (ed) *Qualitative research: Theory, method and practice* Sage, London (1997) pp 8–23

Table 1 Characteristics of a prototypical case of ethnography (the order in which these have been 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 of its various manifestations such that data are gathered from a wide range of sources including interviews, team discussions, incidental conversations, documents, and 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 adopts 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 effects upon the behaviour of the observees.
7. Self-reflection	The observer acknowledges that any interpretative act is influenced by the tradition to which they themselves belong.
8. Intensity	Observations are intensive and long-term so as to enable the observer to become immersed in the ongoing culture of the observee's environment.
9. Independence	The observer aims not to be constrained by any pre-determined goal-set, mind-set or theory.
10. Historicism	The observer aims to connect observations to a backdrop of historical and cultural contingencies.

ventions to traditional ethnography: Bucciarelli's¹⁷ 'ethnographic perspective on engineering design' and Lloyd and Deasley's⁴ 'ethnographic description of design networks'.

Bucciarelli's research is based around studies of two engineering firms. One company was involved in making photovoltaic modules for the conversion of solar radiation into electricity, and the other was engaged in the production of X-ray equipment for a variety of purposes including medical diagnostics, the quality control of materials production, and baggage inspection. Data collection was exclusively by means of field notes. Comparing Bucciarelli's research with the characteristics of a traditional ethnography illuminates a number of interesting counterpoints. To begin with, Bucciarelli gives no indication of the duration of his two studies—although one gets the impression that neither were especially lengthy. Indeed the most detailed aspect of his data which he discusses (relating to the nature of design discourse) pertain to observations deriving from three project meetings. As such, the research seems to violate the intensity of a traditional ethnography. In addition, whilst Bucciarelli claims a high degree of openness in his data collection—which is indeed marked by a strong inclusiveness in terms of the breadth of players who are viewed as relevant

17 Bucciarelli, L. L. 'An ethnographic perspective on engineering design' *Design Studies* Vol 9 (1988) pp 159–168

to the design process and the variety of activities observed—it is evident that he has entered his ethnography with some very strong pre-conceptions of what design *is* and what it is *not*. For example, he is up-front with the views that “designing is more than a cognitive process... Design, even though done by engineers is not a mechanical process... design is not rational in the managerial or economist’s sense” (pp 5–6). Instead, in entering his ethnographic analysis with the clear mind-set ‘that design is a social process’ (p 6) he appears to be explicitly contravening the ethnographic principal of independence from prior theory. Finally, there is no obvious evidence to suggest that Bucciarelli achieved the true status of a participant observer within the design process, which questions the quality of his locatedness within the community of practitioners he studied as well as the mutuality he achieved with his observees.

In spite of its apparent divergence from a traditional form of ethnography, Bucciarelli’s ethnography of design still allowed for the construction of a vivid, multi-layered picture of the commercial design process which is illuminated by the identification of three intriguing themes. First, the concept of ‘object worlds’ is invoked to describe the different design spaces (mechanical, geometric, electromagnetic, managerial etc.) and associated symbol systems that are allied to the various stakeholders in the design situation. Second, the concept of ‘specifications and constraints’ is used to discuss the types of facts and laws that rule over the designer’s practices and limit their development of design solutions. Third, the notion of design discourse being developed into an account of the importance of negotiation from different viewpoints during the design and production of artefacts.

The theme of negotiation is also of paramount importance to Lloyd and Deasley⁴ in relation to their three-month ethnographic study of social processes in an aerospace design team pursuing a 10-year project. Data collection was by means of field notes (since Ministry of Defence regulations prevented audio recording) and involved interviews, analysis of formal documentation, and direct observations of the activity of the electronic design engineers (this latter phase of activity absorbed three weeks of the full study period). Like Bucciarelli’s study, the intensity of Lloyd and Deasley’s ethnography is rather less than that of a traditional ethnography. In addition, the reality of their participant observer status is questionable, and their *a priori* focus on understanding social networks of negotiation in design teams means that their study was in breach of a full degree of openness and independence. Again, though, the value of the work is readily manifest in the interesting findings that it produced. For example, there was good evidence that informal social structures were determining the effectiveness of work activity. In addition, the use of subtle role-playing

in problem-solving was identified as an important aspect of the design process.

The only ethnographic study of design that we have come across which does not violate the feature of a prototypical ethnography is Woolgar's¹⁸ company-based exploration of the processes underpinning the design and production of a new range of micro-computers, which is robust on all ten features and reads as a true ethnography in terms of its narrative richness, socio-cultural locatedness, and epistemological relativism. In our view, the observation that most ethnographies of design diverge from a central case is not something which should be seen as making the research deficient in some way. Indeed, we would argue that systematic differences from the feature list shown in Table 1 are both a necessity and a strength of ethnographic studies of design for both practical and theoretical reasons. To support this view, we invoke the idea that the principal factor that motivates ethnographies of design to differ from a central prototype is that they are *purposive* in ways that traditional ethnographies are not. In particular, almost all ethnographic studies of design that we are familiar with seem to be motivated towards the pursuit of applied goals such as those of providing support to aspects of design cognition (e.g., through technology-based assistance) or improvements to the social milieu of design teams (e.g., through organisational changes and interventions).

2 *The characteristics of applied ethnography*

The way in which such applied goals lead to the violation of the characteristics of a prototypical ethnography are subtle and multi-faceted, and we therefore only summarise three key contraventions here to illustrate our basic argument.

First, the often extreme *intensity* of traditional ethnographic data collection is unlikely to be cost-effective—and may even be impossible—for most design-research projects. Pursuing applied objectives typically brings with it the need to understand skilled behaviours within short time frames in order to effect improvements as quickly as possible. This factor, coupled with the inevitable limitations that are placed on accessing designers owing to management perceptions that work is being interrupted, means that observational intensity will tend to be diluted in favour of temporal specificity, whereby the researcher samples snap-shots of ongoing behaviour over short-term company visits of days, weeks, or at most a few months.

Second, the high degree of *independence* of traditional ethnographic data acquisition and analysis seems inappropriate in applied contexts. Whilst

18 Woolgar, S 'Configuring the user: The case of usability trials' in J Law (ed) *A sociology of monsters: Essays on power, technology and domination* Routledge, London (1995) pp 57–99

traditional ethnographies strive to remain unencumbered by existing theoretical frameworks and classificatory systems, applied ethnographies often require a strong element of hypothesis testing and theory-driven observation. For example, the specification of a new technological system or the revision of existing work practices are both endeavours that come laden with (often implicit) hypotheses that require testing. Note that the need for applied ethnographies to be at least partly theory driven does not exclude some role for genuine *openness* in their research. Indeed, the applied ethnographer's sensitivity to adventitious findings or unexpected observations is a valuable resource for generating novel research hypotheses. In this way, applied ethnography retains a degree of independence which makes it distinct from many other forms of empirical analysis.

Third, the *personalisation* associated with traditional ethnographies, which typically revolve around first-person story-telling in relation to people's historical and cultural endowments, presents a potentially severe problem in applied research owing to the clear lack of verifiability of the outcomes of the analysis. It is important to note, for example, that systems development work in design contexts is inevitably expensive for sponsors, who will quite legitimately be keen to know that the findings deriving from empirical analyses are firmly grounded in relatively objective data. This strong emphasis on verifiability is central to an applied ethnography that can allow for the replication of findings by other observers as well as the validation of results through a process of methodological triangulation.

One key characteristic of this concept of applied ethnography that we have developed so far is that it has no specific allegiance to any overarching epistemological framework for interpreting and explaining human behaviours and practices. It is possible, however, to add such explanatory frameworks (e.g., those deriving from cognitive science or social psychology) to the concept of applied ethnography in such a way as to produce fields of enquiry that fall within the remit of, for example, an 'applied cognitive ethnography' or an 'applied social ethnography'. Indeed, much of our own ethnographic research on design, whilst being motivated toward applied goals such as the computer-based support of information storage and retrieval, is also firmly located within a cognitive-science tradition. As such, the focus of our research is typically that of understanding aspects of process, structure and representation that are inherently *mental* in nature.

Whilst the interpretative frame of reference for an applied cognitive ethnography would relate to the fine-grained explanation of cognitive phenomena with recourse to information-processing constructs, that for an applied social ethnography might well relate to the explanation of phenomena in

terms of processes of social construction. Taking these arguments a step further, we see no reason why a cognitive scientist could not use ethnographic methods in non-applied research aimed at the advancement of a basic understanding of cognitive phenomena, thereby simply undertaking a 'cognitive ethnography'. In making these observations, however, we are not aiming to promulgate a whole set of new terminology for describing different forms of ethnography but are rather attempting to illustrate the ways in which ethnographies of design may not only diverge from a prototypical ethnography but may also exhibit distinctive *flavours* because of differences of purpose and theoretical grounding. Our key argument is that it is typically the applied goals of design ethnographies that drive their divergence from a prototypical case.

3 *The use of applied ethnography in specifying a design support tool*

In this section we illustrate the value of applying ethnography in design research with reference to a three-year research project that we are currently pursuing which is aimed at supporting design 're-use' in innovative engineering contexts through intelligent, computer-based support. Re-use refers to the process by which previously encountered design information (e.g., that relating to performance criteria, solution options, evaluative critiques and the like) is re-invoked by designers when they are tackling a current problem. Existing empirical evidence¹⁹⁻²¹ suggests that previous design concepts *do* frequently get re-used in relation to *routine* design tasks that are characterised by well-defined design problems where known solutions can be re-applied. To date, however, little is known about the nature and extent of information re-use in more *innovative* design situations of the kind that prevail in the conceptual phases of much design work. Designers often tell us that re-use plays only a minimal role in their innovative design work. However, it remains possible that the adaptation of previous design information could make cost-effective use of design effort during innovative design activity (e.g., through the encoding and later retrieval of innovative design ideas that were not technologically feasible at time of generation but that gain applicability over time).

Such inadequacies in our understanding of re-use activities in innovative design contexts make it difficult to know how best to go about facilitating or enhancing information re-use through technological support. This lack of knowledge formed the starting point for the development of the computer-based system for indexing re-use information which we commenced two years ago. This system, called DESPERADO (DESIGN Process Encoding and Retrieval by Agent Designated Objects²²), has involved two parallel research strands. First, on the empirical side, we have been using ethno-

19 Gero, J S 'Design prototypes: A knowledge representation schema for design' *AI Magazine* Vol 11 (1990) pp 26-36

20 Maiden, N and Sutcliffe, A 'Analogously based reusability' *Behaviour and Information Technology* Vol 11 (1993) pp 79-98

21 Visser, W 'Use of episodic knowledge and information in design problem solving' in **N Cross, H Christiaans and K Dorst** *Analysing design activity* Wiley, Chichester (1996) pp 271-289

22 Ormerod, T C, Mariani, J, Spiers, G, Ball, L J and Maskill, L 'Promoting innovative design through indexing agents' Unpublished grant proposal, ESRC Cognitive Engineering Initiative Grant No L127251027 (1995)

graphic methods to investigate the re-use practices of design teams working on conceptual design tasks in large scale commercial design settings. Second, on the computational side, we have been applying ideas from object-oriented programming in the development of DESPERADO's underlying functionality.

The inherently applied objectives of our empirical work readily illustrate the way in which our ethnography had to deviate from a pure form of traditional ethnography. For example, since we knew that we would have to begin the specification and prototyping of DESPERADO in the second year of the project it was quite clear that the intensity of a traditional ethnography would have to be sacrificed in favour of a high degree of observational specificity. In addition, there were two key factors which impacted upon the need for our ethnography to have a strong element of verifiability. First, the fact that DESPERADO had to meet the needs of end-users meant that we had to develop a well-specified and testable account of design re-use that captured both generalities across different companies and contexts as well as team-specific practices within such companies and contexts. Second, we needed to be able to assure our sponsors—and the intended users of DESPERADO—that we had robust theoretical grounds on which to develop the system. Finally, the purposive nature our data collection (i.e., the goal of the empirical work was to elicit requirements for the development of DESPERADO) meant that it was necessary to have some notion of the *kinds* of data that we needed to obtain prior to the actual commencement of our studies. Specifying the type of data that one is aiming to collect in an *a priori* manner clearly compromises the independence of an ethnography, and, to some extent, its openness and richness.

There were, in fact, three particular considerations that dictated the scope of our ethnographic data collection and analysis. First, we needed to identify the *primitive units* of re-use in terms of the basic kinds of information that get re-invoked and adapted during design. Such an understanding was an essential aspect of the requirements specification for DESPERADO since the system would need to base its functioning around these primitive elements of re-use—or what we subsequently refer to as design 'episodes'. Second, we needed to explore the *process* of natural re-use in order to gain insights into designers' strategies for encoding information for subsequent re-use and for retrieving information for current re-use. This knowledge was important for informing the specification of DESPERADO's system-initiated encoding and retrieval functioning. Third, we needed to acquire an understanding of the *situated* nature of re-use practices such as the antecedents and sequelae of re-use. In this respect, evidence for sub-optimal re-use practices (e.g., tendencies to fix rigidly on single solution

options) would clarify the way in which DESPERADO might operate so as to optimise design activity.

3.1 Collecting the data to inform system development

In order to collect data that addressed our themes of interest we conducted short ethnographic studies in our four end user companies— which were all major international leaders in the respective fields of telecommunications, computer systems, industrial design and aerospace. Because of space considerations we will limit our discussion in this section to the study undertaken in the aerospace company, since the data from this study revealed some of the best illustrative examples of generic findings that informed the design specification for DESPERADO.

Prior to the study we held several short meetings over a six-month period with the managers and selected members of the design team. This phase enabled us to set up the study and to negotiate when, and under which terms, our researcher could observe the design group. In addition, these meetings proved extremely useful in facilitating the early capture of basic background information that allowed us to gain familiarity with the design domain of the end-user group.

Our researcher spent a consecutive two-week period with the aerospace design group in the role of a participant observer, in that whilst she did not contribute directly to the design work she was frequently involved in design discussions—often offering opinions and informal assessments of design ideas. The researcher was given a desk in the main design room of the team which was a large open-plan office with limited physical divisions between desks. She was allowed to observe the work of individual designers and was permitted to sit in on team meetings and design-review discussions. The researcher also conducted several informal and unstructured interviews with design team members on a primarily one-to-one basis. Data collection consisted exclusively of field notes that detailed the activities and behaviours observed by the researcher as well as the key content of verbal communications. These observational data were supplemented with notes detailing the researcher's comments upon, and feelings towards, the events and situations that were experienced.

3.2 Using the data to inform the development of DESPERADO

As stated above, we set out to address three key issues; namely, the primitive units of design re-use, the process of re-use, and situating re-use in ongoing design practices. Our ethnographic observations have had pro-

found effects upon the design of DESPERADO in each of these areas, which we briefly describe below.

In attempting to specify the nature and size of re-use primitives our design of DESPERADO was initially guided by the Questions, Options and Criteria (QOC) formalism for capturing design rationale developed by MacLean *et al.*²³. *Questions* describe key design issues, *Options* describe possible solutions to design problems, and *Criteria* describe approaches to assessing options. However, it rapidly transpired from inspection of the ethnographic transcripts that designers do not segment questions, criteria and options in ongoing work. Instead, they are worked upon in parallel, in what we have described as a *focus constellation*. This is in some respects fortuitous, since shifts in question focus, which signal the boundaries of a focus constellation, appear from the data to be a seemingly natural transition point between episodes. Thus, we took the data from the aerospace company and sectioned it into 50 episodes that we have subsequently used to 'seed' the DESPERADO database for company trials of the system. These episodes, with few exceptions, offer re-use information based around the pursuit of a focus constellation by one or more designers at a point in time. Segmenting design episodes by focus constellation affords considerable advantages since DESPERADO is now able to support structured, goal-oriented shifts in activity which we have elsewhere argued²⁴ are a hallmark of design expertise. Furthermore, we make explicit the elicitation of design rationale²³ through the encoding of a focus constellation.

Analyses of the ethnographic data have produced insights into the processes underlying design re-use. Most strikingly, we have seen little evidence of *systematic* encoding or retrieval practices based around external repositories, though there is much evidence of unsystematic storage (and consequent failure to retrieve). Instead, information re-use arises primarily from individuals recalling information from memory. This observation has led us to alter our views on encoding in DESPERADO. Our initial conception, of encoding as a task to be removed from designers where possible, has been replaced by a view of encoding as integral to delivering both *indexing* for re-use and the construction of a design rationale. Thus, instead of attempting to embody inappropriate natural encoding practices, we have refined the DESPERADO interface to make explicit the designer's encoding role, and to support encoding through interface features (e.g., system prompts, defaults and procedurally-oriented displays). This change is justifiable on the grounds that, whilst designers may be reluctant to undertake encoding activity in an unsupported environment where the purpose of encoding is not explicit, the added value gained from achieving both

23 MacLean, A, Young, R M, Bellotti, V M E and Moran, T P 'Questions, Options, and Criteria: Elements of design space analysis' *Human-Computer Interaction* Vol 6 (1991) pp 201-250

24 Ball, L J and Ormerod, T C 'Structured and opportunistic processing in design: A critical discussion' *International Journal of Human-Computer Studies* Vol 43 (1995) pp 131-151

indexing *and* design rationale is sufficiently compelling to undertake the task within a supportive environment.

Perhaps the most striking influence of our ethnographic analysis was upon our understanding of situating re-use within design practices. Prior to our ethnographic studies, we had conducted a large body of studies using the traditional protocol analytic methods of cognitive psychology, to explore the nature of design expertise^{25,26}. A general conclusion from these studies was that even highly skilled designers, when working individually on both small- and large-scale design problems, can be seen to exhibit 'satisficing' behaviours. However, fine-grained analyses of our ethnographic data have revealed unexpected re-use strategies in *team* design. The results are particularly striking in that they contrast so markedly with our previous studies of *individual* design activity. In particular, specific team members (e.g., the project manager) bring to bear strategies aimed directly at avoiding premature commitment to single solution options. This finding has led to a reconceptualisation of the role of DESPERADO. Our previous view of the system acting as a passive 'servant' in support of the tasks of encoding and retrieval of design re-use information has been replaced by a view of DESPERADO as an active 'surrogate manager' that oversees episode traversal during unmanaged design activity, and prompts designers to retrieve previous episodes when there is a danger of satisficing behaviour emerging in individual design work.

4 Using applied ethnography in supporting naturalistic experimentation

So far in this paper we have discussed how ethnography can be modified in pursuit of purposive goals that underlie design activity. Whilst applied ethnography differs from a radial category of pure ethnography, both approaches remain faithful to a non-interventionist research tradition. Most cognitive psychology, on the other hand, is of an interventionist nature, gaining its strength from the systematic manipulation and control of variables that might affect performance. Examples of the successful use of an interventionist experimental methodology to explore natural practices and expertise in design are few and far between. This is hardly surprising given the complexity of the domain, and the obvious difficulties associated with attempting the rigorous control of the majority of factors that might affect design in pursuit of the exploration of a limited subset. Attempts to use experimental methods to explore design immediately confront the problem of *ecological validity*²⁷; put succinctly, the methods used to study performance in any domain must maintain a sufficient degree of authenticity that the researcher can have confidence in the generalisation of their findings beyond the controlled conditions of the laboratory.

25 Ball, L J, Evans, J S B T, Dennis, I and Ormerod, T C 'Problem-solving strategies and expertise in engineering design' *Thinking and Reasoning* Vol 3 (1997) pp 247–270

26 Ball, L J, Maskill, L and Ormerod, T C 'Satisficing in engineering design: Causes, consequences and implications for design support' *Journal of Automation in Construction* Vol 7 (1998) pp 213–227

27 Neisser, U *Cognition and reality* Freeman, San Francisco (1976)

In the face of the problem of ecological validity, one might question whether controlled experimentation has any role to play in design research. We suggest that it does. This role is perhaps best illustrated in the use of controlled experimentation for functionality and usability testing of products and prototypes²⁸. The big pay off from controlled experimentation is that it offers a high degree of objectivity in observation and interpretation. When one needs assurance (as in the case of testing whether a product, component or process satisfies a requirements specification), then controlled experimentation meets this need. Similar requirements for objective assurance face the researcher in developing theories of design expertise, accounts of optimal design practice, or tools and methods for supporting design.

Although ethnographic methods provide rich and naturalistic data for ecologically valid research, they face the inevitable criticism that observations are open to subjective sampling and interpretation. Ethnographers might raise a number of arguments against this criticism, for example by pointing to the skilled nature of ethnography, the potential for achieving validation through intensive observation, and possibly even the false assumptions that underlie claims of objectivity in controlled experimentation. Rather than try to defend ethnography from this criticism, we see a different resolution to the problem of introducing objectivity into design research. We suggest that this can best be achieved through a process of *triangulation* across research methodologies. Triangulation, that is, the support of an observation or interpretation by providing converging evidence from a number of different sources, has been suggested as an important component of the ethnographic method. However, typically the different sources of evidence come from within the same ethnographic study and rely upon the same methodology for their capture. To achieve an acceptable degree of research objectivity, we argue that triangulation must be based upon convergence across different research methodologies. This is made possible in our research because applied ethnography plays a unique role in *enhancing* the ecological validity of controlled experimentation in design research. Thus, outcomes from studies involving ethnographic and laboratory-based methods can be triangulated whilst maintaining ecological validity of the research.

28 Lansdale, M W and Ormerod, T C *Understanding interfaces: A handbook of human-computer interaction* Academic Press, London (1994)
29 Hoffman, R R, Shadbolt, N R, Burton, A M and Klein, G 'Eliciting knowledge from experts: A methodological analysis' *Organizational Behavior and Human Decision Processes* Vol 62 (1995) pp 129–158

4.1 *Sorting—an experimental method that benefits from ethnography*

One experimental technique that has been used to some effect in exploring design expertise, but that suffers from problems of low ecological validity, is the sort method²⁹. In this method, participants are given sets of cards that contain descriptions pertaining to different aspects of a domain, and

are required to sort these into categories according to one or more dimensions that appear important or significant to the participant. The method originates in experimental research into human memory and the mental representation of conceptual categories³⁰. The rationale of the method is that, by examining the nature of the sorts produced by participants (notably the dimensions used to define categories, the assignment of domain objects to different categories, and the order in which category dimensions are produced), one can infer participants' mental representations of conceptual knowledge. In applied research, one typically makes this inference on the basis of studying expert/novice differences in sorting performance. For example, Davies *et al.*³¹ used a sort method to explore the naturalness of different computer program descriptions. They found that sorts produced by expert programmers reflected functional rather than object-oriented groups, a finding that runs counter to the assumed naturalness of object-oriented descriptions of programming concepts.

It has been suggested by some authors³² that the sort method provides an equally informative but more cost-effective method for eliciting accounts of expertise, for example in the development of expert systems, than traditional observational methods such as the analysis of verbal protocols. However, the sort method is not without its problems. In much of the early psychological research into conceptual categories, there was an assumption that what was being elicited through a sort method reflected relatively static conceptual representations. This is certainly the assumption that underlies the use of the sort method by applied researchers^{33,34}. However, this view of static conceptual representations has been challenged by Barsalou's³⁵ description of goal-directed categories. In this view, categories and their members are not fixed, but are determined by the task faced by the individual at any one time. If this view is correct, then an assumption that one can use a sort method to elicit static representations of conceptual knowledge about design is unfounded.

One particular problem facing the use of the sort method in design research concerns the nature of the items to be sorted. In traditional sort tasks, the items consist of object descriptions, often varying in familiarity, representativeness, or prototypicality. In applied research, they may contain descriptions of domain objects or terminology (e.g., programs³¹, programming commands and syntax³⁴, photographs of archaeological shards³²). However, it is in the selection of items for sorting that the method can lose ecological validity. For example, if concepts are goal-directed as Barsalou³⁵ suggests, then the sorts that designers produce are likely to be a function of the items they are given to sort. Specifically, designers are likely to treat items that are presented as representative of domain terminology in

30 Rosch, E 'Natural categories' *Cognitive Psychology* Vol 4 (1973) pp 328–350

31 Davies, S P, Gilmore, D J and Green, T R G 'Are objects that important? Effects of expertise and familiarity on classification of object-oriented code' *Human-Computer Interaction* Vol 10 (1995) pp 227–248

32 Burton, A M, Shadbolt, N R, Rugg, G and Hedgecock, A P 'The efficacy of knowledge acquisition techniques: A comparison across domains and levels of expertise' *Knowledge Acquisition* Vol 2 (1990) pp 167–178

33 Adelson, B 'Problem solving and the development of abstract categories in programming languages' *Memory and Cognition* Vol 9 (1981) pp 422–433

34 McKeithen, K B, Reitman, J S, Rueter, H H and Hirtle, S C 'Knowledge organization and skill differences in computer programmers' *Canadian Journal of Psychology* Vol 13 (1981) pp 307–325

35 Barsalou, L W 'Ideas, central tendency, and frequency of instantiation as determinants of graded structure in categories' *Journal of Experimental Psychology, Learning, Memory and Cognition* Vol 11 (1985) pp 629–654

qualitatively different ways than items that contain the information that they themselves have developed or been involved with. To date, we are unaware of any studies employing the sort method that have used items produced by the study participants themselves. This is not surprising given the difficulty of collecting such items. However, as we have argued in the previous section, applied ethnography can elicit design episodes that provide ecologically valid materials for a study employing the sort method to explore design expertise.

4.2 Combining ethnography and experimentation: an example

We have recently undertaken a study which combined applied ethnography with the use of a sort method. The aim of the study was to investigate whether natural categories that expert designers use to classify design information can be identified. These categories could then form the basis for a notation for encoding design information within DESPERADO. Lack of space precludes a detailed description of the study. However, a brief description of the study illustrates the combination of applied ethnography with controlled experimentation.

The study took place at a major manufacturer of personal computers (PCs), and involved two phases. In the first phase, our researcher spent two weeks conducting an applied ethnographic study of design practices within the company's software development team, which had twelve full-time members. The researcher followed the same approach to that described in the previous section, collecting a mixture of field notes, taped interviews and meetings, texts and other company documentation, whilst noting personal observations and commentary. At the end of this period, in addition to studying descriptions of design practices from the ethnographic observations made by the researcher, we selected from the data a set of thirty design episodes to serve as the basis for the sort study. The episode set was chosen to reflect the different types of design information that were evident in the ethnographic data (see Table 2 for examples).

Participants carried out the sorts individually, in their own office. Each participant reviewed a set of sort instructions that included an example sort. Participants were then given a packet of 30 cards, each card displaying one design episode. Participants were asked to sort the episodes into natural categories segregated by a dimension of their own choosing, and that captured a view of the episodes that was deemed important or relevant to their own design activities. They were told to use as many or as few categories as they found necessary to represent each grouping. Each participant completed two sorts, and the dimensions used to differentiate categories and

Table 2 Examples of episodes used in the software development sort study (note that the example labels were not given to designers, and product and designer names have been changed to protect confidentiality)

Example 1 (audiotaped conversation)

“Actually, ‘Jim’ said to me yesterday, or Friday last week maybe, that ‘John’ wanted to talk about a way of formalising the process from Lotus Notes to NBP. You know, how do we link these and have an actual cut off that says...here’s the point where we stop talking about this stuff and say that’s the decision. Now put it to a Word Pro document and up to the library, and make that a design point.”

Example 2 (email memorandum)

Subject: TARDIS Web Page Fulfilment and Logic Diagram

As the Web page idea has crystallised and grown!, so has the need to firm up the design. I think the logic diagram I put on the [...] Web site [..address..] will describe where we are headed. You will see how the Web page will navigate the user to the right download site regardless of the country or the language. It will also give him the right code that will automatically route his support request to the correct Helpcenter. The response has been no less than enthusiastic from all and we want to keep up the momentum. Comments are appreciated.

Example 3 (field notes)

‘Bill’ describes to ‘Mike’ during an informal morning meeting in Mike’s office how The GUI and ‘engine’ of TARDIS v0.1 are distinct but closely coupled—always running on the same system. TARDIS v0.2 removes this restriction by introducing a new GUI in the form of a combination of HTML web pages and Java applets to allow [..... deleted.....]. Bill is seeking clarification of whether he should continue to work on the new GUI. ‘Mike’ queries whether he is the right person to give this clarification.

Example 4 (Meeting minutes)

TLQ Status 25/02, Minutes—Status Meeting

Awaiting decision from TCP/IP group on bundling BCS with Comms. Level 2 support plan open: Code being tested, problems with ZIP server highlighted and resolved, problems to be resolved with DHCP and TFTP
Actions — Close issues with LBJ (MS < = =person responsible)—LBJ usable version of the code required 25/2 (BB)

the cards assigned to each category were recorded. The resulting data were analysed using cluster analysis and multi-dimensional scaling techniques³⁶, but here we report only a sample of observations.

A number of interesting findings emerge from the analyses. To enable us to detect common patterns in sorting behaviour, the dimensions were summarised and compared for semantic similarity. Table 3 outlines the kinds of dimension produced by the designers. It is notable that, of the 17 sorts generated (one designer was only prepared to undertake a single sort), 15 are captured by just four dimensions. What is interesting is how some of these dimensions, and more specifically the categories generated within them, corroborate the observations made by Bucciarelli¹⁷ in his ethnographic analysis of design. Notably, the dimension ‘relevant to specific group/role’ corresponds with his notion of object worlds, the dimension ‘design process/stage’ corresponds with his notion of the importance of specification and constraint, and the dimension ‘episode content type’ corresponds to his notion of the role played by different types of negotiation.

36 Ormerod, T C, Ball, L J and Rummer, R ‘Verifying accounts of expert design knowledge using a naturalistic card-sort technique’ (in preparation)

Table 3 Dimensions produced by designers in the sort study, with examples of categories generated by a designer under each dimension (note that the labels represent our summaries of the semantic content of dimensions as described by the designers)

Dimension	Designers (d1–d9)
Relevant to specific group/role (e.g., “technical group; systems; bureaucratic; test group; methods; status meeting”)	d1, d2, d5, d7, d9
Design process/stage (e.g., “process; marketing; function; logistics; program development”)	d2, d3, d6, d9
Episode content type (e.g., “obvious procedures; press release; discussion of specification; general discussion; specification details; general narratives; long-winded info of unknown quality; status notes”)	d1, d4, d6
Product (e.g., “TARDIS v0.1; POLAXv0.2; not product-specific”)	d3, d8 (sorts 1 and 2)
Information utility (e.g., “trivial items; info relevant only to participants; sensible suggestions; important work items; hi-level design descriptions; detailed design info that does not affect designers”)	d5
Decision stage (e.g., “Things already done; defined options requiring decision; early design discussions; oddballs”)	d7

Thus, the sort study is able to verify the observations that emerge from this ethnographic study.

The average number of categories produced in each sort was 5.8 (standard deviation = 2.3). Much of the variability comes from a single designer, who produced 10 categories in his first sort, and 11 categories in his second sort. Examples of categories produced are shown in Table 3. In the absence of presenting full statistical analyses, it is inappropriate to comment upon the clustering of cards in each category across designers and consistency of designer’s sorts beyond a few broad observations. However, it is apparent that designers were highly consistent in their category assignments under the dimensions of ‘relevant to specific group/role’ and ‘product’, but much less so under more evaluative dimensions such as ‘episode content type’ and ‘design process/stage’. This is not surprising when one considers the nature of these dimensions, the former delineating category membership on the basis of recognisable function, component or product names, the latter influencing category membership by factors such as the perceived usefulness, importance and priority of design information.

What is particularly clear is that the role played by each designer in the software development team had a strong influence on their categorisation performance. Three of the designers played a senior or managerial role in the team, and their sorts had rather different priorities than the programmers in the team. For instance, the programmers tended to produce more marked or ‘critical’ evaluative categories (e.g., “only relevant to specific

participants”; “trivial information”; “too detailed to be worth inspecting closely”; “excellent program structure”), the managers tended to produce more open evaluative categories (e.g., “detailed decision-making”; “high-level discussion”; “external influences”; “common building blocks”). Differences in sort performance between the various specialist roles within the design team, and the implication that the same design information is evaluated under different metrics across roles, corroborate our observations from the ethnographic study that design re-use may differ in individual and team contexts.

5 *Conclusions*

In this paper we have discussed how the method of ethnography can, and perhaps even should, be adapted for the study of realistic design activities. Applied ethnography differs from a prototypical case of pure ethnography in three principal ways that are associated with its inherent purposiveness: through differences in intensity of observation, in being less independent of prior theory, and in requiring a degree of verifiability or objectivity in interpretation. We have illustrated how using an applied ethnographic technique enabled us to gain important insights into the nature of team design practices that were not possible with more conventional protocol analysis techniques. We were able to sample sufficient data over only a two-week period to make a major contribution to revising the specification of a design re-use support tool in some important ways. We have also shown how combining the outcomes of applied ethnography with a traditional cognitive method such as the card sort enabled us to verify in a more objective fashion the observations made in previous ethnographic studies.

We argue that it is imperative that the differences between applied ethnography, as exemplified in the examples above and practised by other design researchers, and a traditional form of pure ethnography, as espoused by researchers from an anthropological tradition, be clearly stated. Design research that adopts aspects of ethnography without clarifying the exact nature of the adoption runs a number of risks. For example, researchers may find that ethnography falls down on grounds of feasibility and cost-effectiveness. Also, they may fail to recognise the importance of achieving a high degree of verifiability and objectivity where ethnography is conducted in pursuit of meeting design goals and objectives. Furthermore, a failure to clarify the precise nature of applied ethnography may mean that design researchers are faced with what may turn out in the end to be unnecessary and uninformative critiques based upon the assumptions of pure ethnography.

Acknowledgments

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, Gary Spiers and Ryan Rummer 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.