

Human–Computer Interaction: a stable discipline, a nascent science, and the growth of the long tail

Alan Dix

Lancaster University

Abstract

This paper represents a personal view of the state of HCI as a design discipline and as a scientific discipline, and how this is changing in the face of new technological and social situations. Going back 20 years a frequent topic of discussion was whether HCI was a 'discipline'. It is unclear whether this was ever a fruitful topic, but academic disciplines are effectively about academic communities and there is ample evidence of the long-term stability of the international HCI/CHI community. However, as in computer 'science', the central scientific core of HCI is perhaps still unclear; for example, a strength of HCI is the closeness between theory and practice, but the corresponding danger is that the two are often confused. The paper focuses particularly on the challenge of methodological thinking in HCI, especially as the technological and social context of HCI rapidly changes. This is set alongside two other challenges: the development of reliable knowledge in HCI and the clear understanding of interlinked human roles within the discipline. As a case study of the need for methodological thinking, the paper considers the use of single person studies in research and design. These are likely to be particularly valuable as we move from a small number of applications used by many people to a 'long tail' where large numbers of applications are used by small numbers of people. This change calls for different practical design strategies; focusing on the peak experience of a few rather than acceptable performance for many. Moving back to the broader picture, as we see more diversity both in terms of types of systems and kinds of concerns, this may also be an opportunity to reflect on what is core across these; potential fragmentation becoming a locus to understand more clearly what defines HCI, not just for the things we see now, but for the future that we cannot see.

Keywords: HCI discipline, methodology, theory, peak experience, single person study

1 Overview

This paper has its roots in the Inaugural Lecture of SIGCHI Ireland [1], where it seemed a suitable occasion for a sort of 'state of the nation', giving a personal view of where HCI stands as a discipline and how it can develop and grow. This then seemed a suitable topic to build upon for this John Long *Festschrift* special edition of *Interacting with Computers*, as Long has himself written with such sharp insight on the directions of the field and raised questions that have prompted many others to look at the discipline of HCI as a whole, not merely their own work within it. The SIGCHI Ireland talk also drew upon the discussions at the UCLIC/Equator two-day workshop on The Future of HCI in the UK in 2007 (see also [2]), especially the discussion on roles and genres in HCI.

The basic argument of this paper is that while HCI has matured as a community and also as a practice, it is still developing, or needs to develop) its own roots as an academic discipline. I focus especially on methodological thinking, not in the sense of attempting to establish a single methodological framework, which seems doomed to failure, but in encouraging an ongoing methodological critique of the methods we borrow from other

disciplines as we adapt them to our own. This discussion of methodology is set alongside two other, interconnected, challenges to the discipline, the need to establish reliable, validated knowledge (another theme of Long's, although here used more broadly), and the need to understand the way different roles within the HCI community fit together and should communicate their results in order to produce a stronger science as a whole.

Note that I will occasionally use the term 'science', but in a broader, perhaps more historically felicitous sense, than the 'hard' sciences. Much of HCI research (but not necessarily all, see section 3) is oriented towards producing knowledge that is usable for design, hence that has, in a sense, 'truth' about the world. This knowledge may not be quantitative, or formal, but does need to be what I term *synthetic* [3]; that is helping you to achieve some effect, or, in Rogers' [4] terms, *prescriptive*, "providing advice as to how to design or evaluate". In the broad academic dichotomy this seems the domain of the Sciences as opposed to the Arts.

As this is a long paper, Figure 1 shows a roadmap of the main argument structure. The next section provides a motivation for this work, noting the strengths of HCI as a growing discipline and the need to develop stronger independent disciplinary roots. Section 3 takes a short detour and discusses the issue of what HCI is about, its ultimate purpose and goal as a discipline (as opposed to the purpose of individual pieces of work within the discipline); this is largely to contextualise the work within that of others, most notably Long, who have focused strongly on this question. In contrast this paper is more about how the discipline is conducted, what it *does*. Section 4 picks up this theme proposing and discussing three challenges for the discipline: *methodology*, *knowledge* and *roles*. The need for clarity in understanding methodology is particularly important given the rapid changes in the technology and human practices that are the subject of HCI. Section 5 discusses these changes, probably familiar ground for many readers. This then leads, in section 6, to a case study of methodological thinking in the work of Razak on *single person studies*. This technique is of particular use when applied to design for *peak experience*, best for some as opposed to good for all. This trend towards the *long tail* of more individual micro-applications and the democratisation of information technology seems to be common to many of the developments in the use of technology and presents its own challenges for traditional HCI practice.

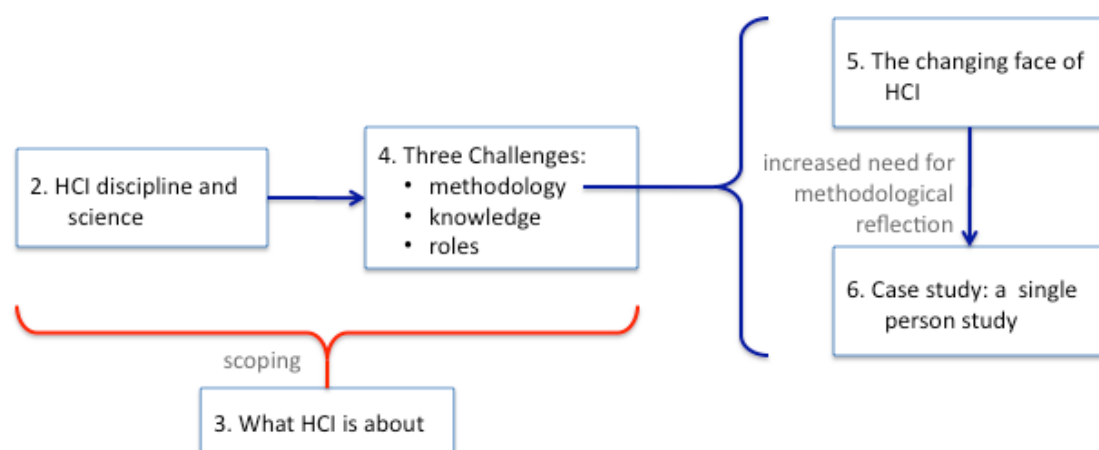


Figure 1. Roadmap of the argument structure of the paper

The paper devotes a substantial part of section 4 to discussing evaluation and the whole of section 6 to the use of single person studies. In both cases these are used as examples of the importance of clear methodological thinking within HCI; although evaluation is

itself a central topic in the discipline. It is hoped that the paper gives insight on both these topics; however, the main focus is on the broader issue of methodology alongside the other core challenges of knowledge and roles in HCI as an integrated community and a scientific discipline.

2 HCI discipline and science

2.1 The Growth of a Community

The roots of HCI go back at least 50 years, with Brian Shackel's paper on ergonomics of displays in 1959 [5]. However, the real beginnings of HCI as an emerging discipline are more like 25 years ago, with the founding of early conferences: Interact, CHI, British HCI and Vienna HCI (now ceased). My own first international conference was Interact'87 in Stuttgart at which Brian Shackel gave a plenary, a welcome, I believe, on behalf of IFIP TC.13¹. A key question he posed was whether HCI was a discipline, or merely a meeting between other disciplines. Now this seems rather like navel gazing, but, at the time when HCI was developing coherence, it was a significant question.

Looking back now we can easily say, "of course it is an academic discipline", because what is an academic discipline if it is not an academic community; after 25 years of IFIP TC.13, and numerous national societies: SIGCHI, Interaction (formerly known as British HCI Group), not least the recent SIGCHI Ireland – clearly there is a community!

But that is a little too glib. Science, using the word in the broadest sense, goes beyond community; to be an academic discipline also requires a coherent basis for knowledge. Mere acceptance of knowledge by a group is not sufficient; we need some assurance of the truth or validity of our knowledge. When flying we are not happy to rely merely on the accepted *opinion* of aerospace engineers; we want to know that they have a basis for their designs *beyond* accepted practice.

2.2 Craft or Science

This brings to mind the discussion in the late 1980s, initiated by Long and Dowell [6], about whether HCI was a craft, engineering or science. Arguably craft is really more about individual experience, but craftsmanship is not what one would want in an aeroplane, nor in Internet banking. Whether we call it science, engineering or simply being academic, we need to be able to give away knowledge to others who should then be able to apply it with assurance. However un-politically correct it is to use this sort of positivist language – yes we do want *truth* and *fact* sometimes!

Note that in talking about the facts of HCI we need to distinguish between the domain studied, which is by its nature complex and nuanced, and our understanding of it, within which we seek clarity. Of course, we rarely, if ever, have complete knowledge – there is always an epistemological gap; and the level of confidence we accept varies from domain to domain. However, for many aspects of HCI research, while the subject matter is often culturally determined and rich, facts about these contexts are not matters of opinion. Similarly, in design itself, while the goals of design and the people for whom the product is being created are contextually and culturally situated, the success of methods in achieving appropriate designs is not.

So, are we getting there; are we developing this *coherent* basis for knowledge in HCI?

¹ IFIP TC.13, International Federation for Information Processing, Technical Committee on Human–Computer Interaction. <http://www.ifi-hci.org/>

2.3 Second Generation HCI

The demographic of the HCI community varies between countries, but certainly in the UK there are an increasing number of 'second generation' HCI people; that is people who have done PhDs, masters and maybe even undergraduate courses with a strong HCI element and have now become the teachers and senior researchers themselves. I would describe myself as originally a mathematician who moved into HCI, others have roots in psychology, computer science, sociology, but an increasing number are straight HCI people.

As a sign of community this is very powerful; no longing for a half-remembered academic homeland elsewhere, but an academic generation who own HCI as their home.

However, this has also given me concern for a number of years. As we gradually lose those strong connections with our old disciplinary roots, have we developed equally strong ones within HCI?

2.4 From Community to Discipline

Indeed, there are signs that we do not yet have strong enough methodological roots. One key example is the relationship between HCI research and practice. One of the great strengths of HCI is that the two are close. To some extent this is true of computing also, but even more so in HCI. There are few fields where the practitioners and researchers can so freely attend the same events, present work to one another and hold discussions. Again this is powerful for HCI as a community and good news for funders looking for industrial relevance; indeed, in many countries, it has been commercial pressure that has driven often reluctant computing departments to take HCI seriously. However, the danger of this is that it is easy to confuse the two. Nowhere is this more evident than in usability evaluation.

We all know that evaluation is the *sine qua non* of HCI; professionally, often the key role of the usability practitioner; and, academically, try getting a paper published without that evaluation box being ticked! The techniques and tools for evaluation are often the same for usability practice and HCI research, whether formal experiments, usability labs, ethnography, prototyping, or maybe even cultural probes or technology probes. However, whilst the techniques are similar, the *goals* are different. For the usability professional, the ultimate aim is to improve the product, whereas the goal of research is to gain new understanding. In fact, even these goals are interlinked: research systems often need to be designed well enough for effective experimentation or deployment; and effective design will be based on a thorough understanding of the context and technology².

However, for the researcher this formative creation of an experimental prototype is NOT the research itself, but merely the preparation for the research; and for the practitioner the understanding they gain is primarily in order to design better systems now, not establish fundamental knowledge for 10 or 20 years' time.

² In the recent Being Human 2020 Microsoft report [7], 'Understanding' is proposed as an important additional stage in the HCI lifecycle; one wonders what designers were supposed to be doing before! To be fair the meaning of 'understanding' in the report about understanding human values, so really the report is talking about an extra thread or concern during the whole evaluation-analysis-redesign part of the design cycle in parallel with other concerns. However, the fact that 'understanding' can, however accidentally, apparently be assumed to be absent from normal interaction design is surely symptomatic of a deep issue to be addressed within HCI.

So here we have a great strength of our community, but one that needs a clear understanding of purpose in order to contribute to a well-founded academic discipline. Reading any conference proceedings or journal it is evident that this clarity of purpose is not yet there. We clearly have work still to do.

3 What HCI is about and how it goes about it

Long and Dowell's [6] seminal paper on whether HCI was craft, science or engineering, and moreover Long's elaboration of the relationship between HCI research and design[8], were predominantly about what HCI is *about* as a discipline: the subject matter, what it produces; the latter paper in particular focused on framing the "HCI general design problem". Similarly Diaper in his opening editorial for *Interacting with Computers* [9], and re-iterated more recently [10], suggests that the goals of the discipline of HCI are "to develop or improve the safety, utility, effectiveness, efficiency and usability of systems that include computers".

Of course, many would object to both Long's and Diaper's apparently Taylorist approaches to the definition of HCI's purpose, certainly adding experience [11], or even human values [7]. However, it should be noted that Diaper and Sanger [10] expressly state that for them 'work' is taken to encompass leisure and is the full activity of being human; that is, satisfaction is subsumed within effectiveness. Similarly Long [8] states that work includes "office work, factory work and home work" (albeit followed by the rider "any activity seeking effective performance"). However, this is all part of the discussion of what HCI is about.

Note that this ultimate 'purpose' of HCI is distinct from the goal of an individual piece of work discussed at the end of the last section. The goal of an evaluation of a prototype within HCI research should always be for gathering understanding not improving the prototype as an artefact. However, the larger purpose of obtaining that understanding may be to help others engaged in practical design to improve their own devices and systems or maybe purely for the sake of the understanding itself.

Although unavoidably based on the author's personal prejudices about the ultimate purpose of HCI, this paper attempts, not to avoid this issue, but to address a slightly different one, namely, how HCI as a discipline goes about addressing its goals, concerns, purpose: that is how it *goes about* doing *what it is about*. Thus the breakdown in the next section is fundamentally about the community of HCI:

- (i) roles – how the community is constituted
- (ii) methodology – how the members of the community act
- (iii) knowledge – how the member of the community communicate

Note that the 'knowledge' here has two purposes, (iii.a) communication *within* the community and (iii.b) communication to (or with) those *outside* the community. As I am focusing on the academic discipline of HCI, 'outside' here will include HCI practitioners (when acting in that role; a practitioner may also be an academic and vice versa). It is the latter, (iii.b), that Long [8] is referring to when he describes the 'discipline' of HCI to be the "use of HCI [knowledge]³ to support HCI practices ...". In contrast, the focus of

³ In the original article the word 'theory' is used instead of 'knowledge', but in a footnote Long explains that the word 'theory' is only used because it was part of a colloquium on

this paper is more on the internal communications (iii.a) that together build a coherent and reliable body of knowledge, but of course the availability and reliability of that body of knowledge is exactly what is needed for its exploitation by those 'outside'.

It may be evident that human interactions, including conferences and discussions, are not included explicitly in my breakdown, although they are of course often the place where 'knowledge' is first presented and also often the place where it is conceived. In fact the breakdown is not precious and it may be extended to include social and community events explicitly in this picture. Indeed, in reporting on the *methods* used in DEPtH, a project focused on understanding physicality in design, we remarked on the central importance of events as a source of knowledge and data as well as a place for community building [12]. That is, the events were important in DEPtH as part of its *methodology*.

So, the focus on externalised knowledge as the locus of community interaction rather than face-to-face meetings may seem odd, but effectively it acknowledges the different ways in which knowledge is passed on. On the one hand there is a form of *diffusion* or *contagion*, where knowledge passes from person to person. This is not to suggest that knowledge is in lump-like memes [13], like passing on a library book; the process is much richer, more one of knowledge being formed and informed in the relations between people. We may discuss concepts with colleagues, at conferences, and in online forums, building up an understanding within the academic community. We may also discuss with practitioners, both learning about their concerns and experience and also passing on the research community's accumulated and distributed understanding.

This human process is of the utmost importance. It can be seen as a form of long-term establishment of *common ground*. While the theory of common ground [14,15] is normally applied to conversations or similar discourse, the same arguments for growing mutual understanding apply to communities albeit with slightly different ramifications and mechanisms. This human process is also a crucial part of education.

However, while the human process of diffusion and mutual understanding is important, it is not sufficient; one of the distinguishing features of an academic discipline is in its *externalisation* of knowledge. We communicate through papers, videos, and other artefacts such as software (although it rarely continues running for long). This is by nature de-contextualised and abstracted, inasmuch as any externalisation abstracts from its subject, but precisely because of this it is applicable and persistent. The published extant knowledge of the discipline is the defining *boundary object*⁴ of the community.

Now in fact this picture is itself idealised, and some might say naïve, in that the interpretation of externalised knowledge is itself governed by understanding gleaned usually through diffusive processes. My own *alma mater* discipline, mathematics, is an arch example of this, where there is a vast, largely unwritten, understanding of processes and interpretation that accompanies the formal mathematical theorems and proofs. However, while it is certainly naïve to ignore the rich community that surrounds

'theory', but 'knowledge' better encompassed his meaning. Here, in the spirit of Hansard, the quote has been modified to reflect his intention rather than his words.

⁴ The notion of a boundary object [16] is normally applied to smaller groups, but applies equally well at the level of a trans-national community. Indeed, Star's original analysis was of taxonomies in a museum environment, not so far from a body of shared HCI literature and resources.

academic externalised 'knowledge', both practical and theoretical concerns suggest this externalisation as an ideal, or at least a touchstone.

As noted, this paper is not primarily focused on the purpose of HCI, what it is about. Still, the dichotomy between inside and outside, academic and practitioner deserves a few words.

Long's own work situates HCI *as an academic discipline* (or in his words HCI research) in relation to HCI design as seen in the quote above or his phrase "*fit-for-design-purpose*" [8]. Similarly Diaper, whilst critiquing the common use of 'usability' on its own as a synonym for HCI, still gives goals that are focused on design, albeit interpreted more broadly than pure usability: "safety, effectiveness, efficiency, and usability". Coming from a different, more qualitative, position within HCI, Rogers also, in her excellent review and critique of (social and psychological) theory in HCI [4], effectively situates HCI theory in relation to its utility in HCI practice.

In teaching I usually distinguish between HCI as an academic discipline and HCI as a design discipline. The latter concerns using skills, knowledge and processes in the production of devices, software and other artefacts (or more generally interventions [17]) that in some way influence human interactions with computers (or more generally technology). The former, HCI as an academic discipline, is the study of situations involving people and technology (note, the series name of the British HCI conference), the design practices involved in such situations, and tools and techniques that are or can be used in either. In order to clarify this distinction, the remainder of this section will use "HCI Practice" to denote the set of situations and design practices and "HCI Research" to denote the academic discipline.

This separation is relatively uncontentious and accords with all the positions described so far; in particular "HCI Practice" here is close to or identical to Long's "HCI design". What is critical is whether HCI Research is simply *about* HCI Practice, or whether HCI Research is *for* HCI Practice (and usually for its improvement). If the latter is the case, then HCI is purely a vocational discipline focused on its practical outcomes⁵. The, now slightly dated, use of the term 'usability' for HCI or, more commonly now, 'interaction design', both orient us to regard HCI as *just* vocational. Now the word 'just' in no way minimises the importance of the vocational use of HCI, nor minimises its importance within HCI, but challenges whether it is the *sole* end of HCI.

The confusion already alluded to between the methods in HCI is partly fuelled by this apparent identification of HCI Research and HCI Practice. It would be consistent to maintain the distinction between HCI Practice itself (the practice) and HCI Research about-but-for-the-purposes-of HCI Practice. However, this muddling of the two in methodology does seem symptomatic of a general lack of distinction.

The roots of HCI undoubtedly began in the practice of creating systems and situations where people could effectively (and all the other adjectives) interact with technology. However, as HCI has grown, it has encompassed a whole study of human endeavour and activity, which appears to go beyond purely vocational aims. Probability was born out of gambling, but does not restrict itself to studies funded by bookmakers.

⁵ Vocational here is being used in the sense common in UK education; that is a subject the primary focus of which is learning in order to practise a trade or profession, or to produce some other practical outcome.

Whether or not one accepts the arguments for pure ‘curiosity-driven’ research, and even if one takes a utilitarian approach, there are good reasons to believe that HCI Research *for* HCI Practice is often best served by focusing on HCI Research *about* HCI Practice. An excessive utility focus tends to mean that research runs behind technology. Work on the newest thing is too late for it, and looking for the next big thing is almost bound to fail; the big win is in using the new, the old and ideas for the next as means of uncovering more general knowledge. It is this knowledge that will be of value for the next new big thing and the next and the next after that. For example, consider the explosion in Facebook research *after* it reached its first 100 million; this has applied some existing theory from *outside* HCI, but appears to be largely a matter of the field catching up on a phenomenon after it has happened, rather than informing the development and design. This clearly happens in any discipline when there are major changes, but it should be the exception not the rule.

Perhaps most worrying is that we can start to accept a technology-driven approach as normal. At CSCW 2002 Stewart Brand gave a keynote taking the idea of timescales of structures in his book "*How Buildings Learn*" [18] and generalising this to look at different timescales in other kinds of activity. He tentatively suggested that there was perhaps a similar set of timescales for research (longest timescale), development (faster) and production (now). However, one of the questioners suggested that in fact in CSCW (and read also HCI) research was if anything faster-moving than development, driven constantly by the latest technology; the audience seemed to agree without noticing the unintended irony of the situation following a talk that emphasised the importance of long-term thinking.

This all said, the rest of the paper stands neutral on this (important) issue of what HCI is ultimately about, and instead focuses more on the way in which we as a community orient ourselves to create reliable knowledge no matter what form it takes.

4 Three Challenges

To recap, HCI research often has trouble distinguishing itself from HCI practice. Whether or not one believes that the purpose of HCI research is to serve HCI practice, clearly there is a difference between the two. However, this seems to stem from a lack of solid methodological roots as HCI has grown away from its parent disciplines, but not clearly established its own methodological heart.

Considering this, I propose three challenges, defined in the previous section, that we need to address in order to develop the academic discipline of HCI: challenges of methodology, of knowledge and of roles. As noted these are primarily challenges addressed to the conduct and process of the community in relation to building a systematic and reliable body of knowledge. I am not addressing the (more thorny) issues of what HCI *is* and what it is *about*. Even within this narrower remit, there are surely many other challenges, not least that of the inter-human relationships noted; however, three is enough to start with and is the classical number of points in any argument.

4.1 Methodology

New disciplinary roots require new methods. We have many methods in our HCI toolkit, so this does not seem to be a problem. However, these methods are often ‘borrowed’ from other disciplines. Within an established discipline one can use accepted methods without a great deal of thought as to whether they are appropriate because one is using them in the same way that others have before. But if we simply adopt these methods in

a new context without considering the desiderata that made them appropriate in the original context, they may be misleading or lead us to conclusions that are downright wrong.

To adopt methods in a new discipline means we have to understand *why* the methods work; that is we have to think *methodologically*. Now here I am not using methodology in the way that has become common in computing and in HCI; that is simply to mean method! When I say we need to think methodologically I mean we need to think *about* our methods.

The frequent confusion between usability evaluation and evaluation in research is just such a failing – adopting methods without understanding methodologically why they are appropriate and what they are for. But to be fair this is because thinking methodologically is hard. Many established disciplines have been around hundreds of years and so have had time to develop or evolve appropriate methods, but if we look to newer disciplines or sub-disciplines we often see methodological problems.

As an example, and at the risk of alienating a community ... some years ago I was evaluating some work that was on the boundaries between HCI and distributed artificial intelligence (DAI). The work seemed to be fundamentally flawed, as the researcher had performed single runs of a stochastic simulation with different conditions. In an HCI setting this might be like running an experiment with a single user in each condition. Now because of the nature of simulations there are times when a single run is sufficient; when certain conditions hold, a long enough run will exhibit all possible behaviours and so one long run is effectively like doing lots of short ones. Unfortunately in this case, the runs were too short due to memory problems. I was worried that I was going to have to send the researcher back to do repeat runs of everything! Happily there was a second reviewer who came from the DAI community and so I was able to ask him about this ... what was the accepted practice in the community. Again happily the other reviewer was not just from the community, but also had a deep *methodological* understanding of the history of the area. Early in the development of DAI a key figure had shown that single runs were acceptable *so long as the relevant conditions applied* ... unfortunately the 'single run' part got remembered, but the conditions had been forgotten. The research we were evaluating was simply following accepted practice in its discipline ... it is just that the accepted practice was flawed.

Now before you judge DAI too harshly, think how often you have read papers (or written them yourself) that quote "you only need to test with five users" [19] or Miller's 7 ± 2 [20] without really checking that they are valid in the context? In the case of the figure of five users, this was developed based on a combination of a mathematical model and empirical results [21]. The figure of five users was calculated

- (i) as the optimal cost/benefit point within an *iterative development cycle*, considerably more users are required for summative evaluation or where there is only the opportunity for a single formative evaluation stage;
- (ii) as an average over a number of projects and to be applied properly needs to be assessed on a system by system basis; and
- (iii) based on a number of assumptions, in particular, independence of faults, that are more reasonable for more fully developed systems than for early prototypes, where one fault may mask another.

Just as in the DAI example, Nielsen and Landaur's original paper outlines many of these limitations, but the finding is typically used without similar care.

Similarly, Miller's 7 ± 2 is about working memory [20], yet is frequently applied incorrectly when, in fact, other cognitive or visual processes are important, for example, the length of menus. In fact a visual menu does not require much working memory so long as it is organised clearly, and in laboratory experiments Larson and Czerwinski [22] found that far broader menus were optimal. One can easily work out simple back-of-the-envelope models of menus to calculate the optimal trade-off between breadth and depth based on time to search the menu visually and time for the page to refresh [23]. For the web, where visual search is typically much faster than refresh time, the optimal figure is typically 60-plus items per menu. This said, Miller's 7 ± 2 may be useful when applied to the *depth* of the menu, if users try to keep track in their head of where they have been. There is some evidence that older users, often with poorer working memory, find deep menu structures particularly difficult [24]. For more examples of misuse, Eisenberg [25] produced an excellent critique of 7 ± 2 from a designer's viewpoint, although he doesn't appear to have realised that the poor uses of it are actually misuses.

As well as misusing past results, it is easy to go fundamentally wrong in the application of empirical methods whilst deriving new results. A vignette I have used several times previously [26,27] illustrates the importance of methodological thinking and what can happen when this fails. It concerns a paper that was published at a major ACM sponsored conference in HCI a few years ago. To spare the authors, I will not name the paper. It may sound familiar; however, this is likely to be because it is typical of many papers that fail methodologically in the same way.

Note that the paper being described was not only in a conference with one of the most rigorous reviewing procedures, but also included among its authors a major figure in the field. The kind of methodological problems described below are not simply those of students or junior researchers, but common amongst the most senior figures in HCI.

The particular paper was in fact a solid empirical paper: experiment, design, and evaluation. It was considering collaborative support for a task; call it task X. The work began by considering three pieces of software; call them A, B and C:

- A domain-specific software, synchronous group interaction
- B generic software, synchronous
- C generic software, asynchronous

The three systems were placed in a 2x2 matrix: domain-specific vs. generic on one axis and synchronous vs. asynchronous on the other (Fig. 2). Incidentally, these 2x2 matrices are ubiquitous in many areas and are indeed extremely powerful analytic tools [28,3].




Figure 2. Systems placed in a 2x2 matrix

The paper then went on to describe the experiment comparing these conditions. There were a reasonable number of subjects in each condition (not just 5!) and sensible quality measures were used for assessing the outcomes of the task. Furthermore the experiment revealed statistically significant results ... well certainly $p < 0.05$. There were two main effects:

- (i) domain-specific software was better than generic software (Fig. 3a), and
- (ii) asynchronous was better than synchronous (Fig. 3b).

The paper then concluded that what was clearly required was the missing gap: domain-specific asynchronous software, and then went on to describe the design and evaluation of an application in this area.




Figure 3. Experimental results (schematic)

This all sounds exemplary, so what's wrong with it?

First of all the paper was a little strong in its suggestion that (i) and (ii) meant that domain-specific asynchronous software would be best of all. Interaction effects⁶ are very common in HCI, and there was no argument as to why we shouldn't expect an interaction in this case. However, that said, certainly the results would suggest that it is a good case to investigate further.

However, the big problem is harder to see, indeed if you blinked at the wrong moment when reading the paper it was easy to miss entirely. The paper started off with three systems and quite properly analysed them along dimensions. However, it then went on to conduct the experiment as if there were two independent variables being manipulated, when in fact there was precisely one piece of software for each condition. This is analogous to having an experiment with precisely one user in each condition – clearly problematic, and yet it is common to use a single piece of software, just as in this case, and never realise there is a problem!

In case it is not obvious why this is so bad: these were three completely different pieces of software that happened to have the relevant properties. Suppose application B just happened to have been very poorly designed. Application B would perform worse than application A, giving rise to the apparent effect that domain-specific was better than generic (Fig 3a). Application B would also perform worse than application C, giving rise

⁶ An interaction effect is when one cannot simply 'add up' the independent effects. For example, just because you like hot fish better than cold fish, and cold ice cream better than cold fish, doesn't mean you prefer hot ice cream overall.

to the apparent effect that asynchronous was better than synchronous (Fig. 3b). That is the effects may have been due to an entirely extraneous factor, and nothing to do with the actual properties being studied.

The problem here is that the paper (and many in HCI) has 'borrowed' controlled experimental methods from psychology, but these methods embody assumptions that often do not hold in HCI. In particular, controlled psychological experiments are designed so there is a single simple cause or manipulation between conditions. However, when used in HCI, as above, there are many uncontrolled causes. Often we want experiments that have some form of ecological validity, which makes this worse.

However, just imagine trying to run the above study as if it were really a heavily controlled psychological experiment: we take a single piece of generic synchronous software B, first tweak it so it becomes domain-specific (call it A) and then tweak it again to make it still generic but asynchronous (call it C). This sounds stronger, and sometimes can work. However, in this case (and many) you would need to take a piece of software that is well designed for a particular situation (task generic and synchronous) then change to 'just' make it domain specific or 'just' make it asynchronous; and furthermore you would need to do this without changing anything else and have it equally 'good' in all other respects after the tweak ... clearly not possible.

Does this mean no experiments are possible in HCI? Far from it; by understanding the assumptions underlying controlled experiments and the way in which HCI experiments do not meet these assumptions, we are in a position to alter the practice of the experiment and the methods of analysis in order to make more reliable interpretations of the outcomes.

In this case, we could collect additional qualitative data (video, logs, audio) as is common in HCI, but then use these in order to help interpret the quantitative measures. Based on our knowledge of human interaction and the data that was collected, the researchers might have been able to come to some *judgement* as to whether the effect seen was to do with the difference between synchronous and asynchronous interaction, or due simply to specific features of application B or C.

Quantitative end-to-end measures are good at telling you *whether* there is an effect, and *how strong* it is, but it is the qualitative data that helps you to understand *why* you are seeing the phenomenon. Furthermore, richer data can help reveal the *mechanisms* underlying the observed phenomena. By 'mechanism' I mean the details of how a person engages in some form of activity or task including, where appropriate, observable phenomena, social and cognitive aspects. That is, a detailed account not just of what happens end-to-end, but the steps, actions and thoughts that are necessary between. When you understand mechanism, it is easier to see whether an overall result is likely to generalise to new situation, and to address empirical or observational data in an analytic manner [3]. Furthermore, if you understand mechanism then you may be able to add new measures or interventions to study finer aspects of the overall interaction.

The importance, and sometimes difficulty, of establishing mechanism is important in other fields also. Nutritional advice for many years quoted a recommended daily protein intake far higher than today. The reasons for this lay many years before when protein deficiency was first recognised during a famine in Africa. Once the medical team studying this realised the disease they were seeing was due to protein deficiency, they began to administer protein supplements until the signs of disease went away and concluded that this was the required amount of protein. What they did not realise was that the body also burns protein for energy if other sources (fats and sugars) are not available. The children involved in the programme were starving and the team were

first providing sufficient calories in the protein supplement before the children could start to use it properly as protein. Like HCI research, nutritional research involves many complex interconnected factors and it is often hard to alter one without affecting others. This makes it even more important to understand the underlying physiological mechanisms, however difficult this may be, in order to prevent major mistakes.

While this section has focused particularly on evaluation and in particular empirical evaluation, it is here to serve as an example of the wider issue that we need to explicitly think about methodology in all its manifestations.

4.2 Knowledge

To be a strong discipline, we need ways of gathering sound knowledge, ways of knowing what is true, and ways of establishing validity. As part of the Future of HCI discussions in the UK in 2007, Blandford [2] emphasises that “HCI research delivers new insights ... that are valid” and validation is critical also in Long’s [8] discussion of the relation between HCI research and design.

As previously noted, evaluation is central within HCI, to the extent that if one wanted to point to a touchstone for what is *accepted by the community* to be the sign of valid research, surely evaluation would be it. The major exception is pure studies of existing work-practices in domains where technology is already present or expected to be introduced. Otherwise whatever kind of ‘thing’ you have produced as an outcome of your research, be it a concept, a method, a toolkit, or an application, what the reviewers of your paper want to see is some level of evaluation and typically evaluation with real users.

Of course seeking some form of validation of your work is critical – after all I have said that we are after truth not mere opinion. However, it is wrong to assume that evaluation is the only means to verify validity. In mathematics, you do not evaluate a theorem to see if it is true, you *prove* it, that is provide a *justification* of why it should be true. Mathematics is unusual in being able to put *all* of its trust in justification; the particular closed nature of mathematical argument makes this possible. In general, academic disciplines vary as to the relative importance of evaluation (Fig. 4); in particular, evaluation is more important where the phenomenon being studied is complex or hard to predict, or ability to reason may be limited. For example, in medicine one might establish, based on theory or prior art, that a particular family of compounds is likely to be effective in treating a condition (justification), but the complexity of the human body and pharmaceutical chemistry means you need laboratory studies and eventual clinical trials to find out which actually works (evaluation).

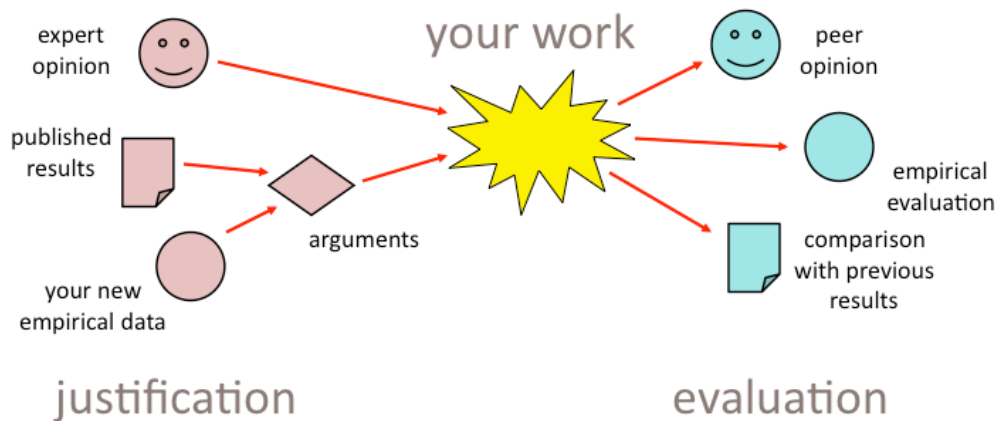


Figure 4. Forms of validation

Arguably if our work is *only* validated through evaluation it is pure invention, not academic research at all – after all we should have some reason for what we do, not just randomly trying what occurred to us in the bath one morning.

Evaluation is especially problematic for *generative artefacts* – that is things that in some way make other things or can be instantiated in different ways [3,29]. This includes theories, methods, guidelines, tools, architectures ... indeed just about anything we produce as research outputs in HCI! The problem is that evaluation cannot exhaust all possible uses or instantiations of a generative artefact, so can never validate it fully. Indeed, as an easy to remember catch phrase (-)):

the evaluation of generative artefacts is methodologically unsound [29]

Even a single piece of software is a generative artefact as it is only in the specific moments of use that it becomes grounded. We cope with this in usability testing by trying to have sufficient users working on a sufficient range of tasks in order to *sample the space* of potential use. However, once we get to design notations or guidelines, sampling becomes all but impossible. To say we have reasonably covered the space we would need to get many different designers with many different briefs and then usability test each outcome with many different users ... and that is just to answer the simple question "does it work?", let alone "why?" and "how can we improve it?".

Note that this is not to suggest that empirical evaluation plays no part in validation of these complex generative artefacts such as methods; it is just that any empirical evaluation needs to be part of a theoretical argument or some other form of justification. As an example of this, Furniss *et al.*'s [30] recent work studying usability evaluation methods is exactly adopting this approach, using a combination of observations of practitioners using different methods, but set within a theoretical framework including distributed cognition and resilience engineering theory. Related work on has shown how various usability evaluation methods have different scope as to what aspects of the design they are best suited [31]; that is each forming part of a larger argument or process.

Within HCI there are a gamut of techniques available for both justification and evaluation, including for justification:

- existing published results of experiments and analysis
- one's empirical data from previous experiments, studies, etc.

- expert opinion (published or otherwise) and common sense
- arguments based on the above

... and for evaluation:

- fresh empirical evaluation, user studies, timing data, etc.
- peer reviews of one's work (do other people agree it is a good idea)
- comparison with previous work (do the parts that should behave the same actually do so)

In any field, the powerful thing is how these *work together* to establish validity. Even in mathematics, the domain of pure justification, it is common to try out a potential theorem against example data either to look for counter examples (Popperian falsification [32] is evaluation based), or to suggest how a proof might proceed: anyone who has done geometry in school will have experienced this using sketched triangles and circles at the beginning of a proof. Here the evaluation is guiding the process of justification. This can be the case in HCI: as you notice patterns in empirical data you think "of course, that must be because ...".

Equally important is that when one builds the justification of why something should work, the argument will *not* be watertight in the way that a mathematical argument can be. The data on which we build our justification has been obtained under particular circumstances that may be different from our own, we may be bringing things together in new ways and making uncertain extrapolations or deductions. Some parts of our argument may be strong and we would be very surprised if actual use showed otherwise, but some parts of the argument may involve more uncertain data, a greater degree of extrapolation or even pure guesswork. These weaker parts of the argument are the ideal candidates for focusing our efforts in evaluation. Why waste effort on the things we know anyway; instead use those precious empirical resources (our own time and that of our participants) to examine the things we understand least well.

This was precisely the approach taken by the designers of the Xerox Star. There were many design decisions, too many to test individually, let alone in combinations. Only when aspects of the design were problematic, or unclear, did they perform targeted user studies. One example of this was the direction of scroll buttons (see Fig. 5): should pressing the 'up' button make the text go up (moving the page), or the text go down (moving the view)? If there were only one interpretation it would not be a problem, but because there was not a clear justification this was one of the places where the Star team did empirical evaluation ... it is a pity that the *wrong* answer was used in subsequent Lisa design and carried forward to this day, but that is a different story! [33,34,35]

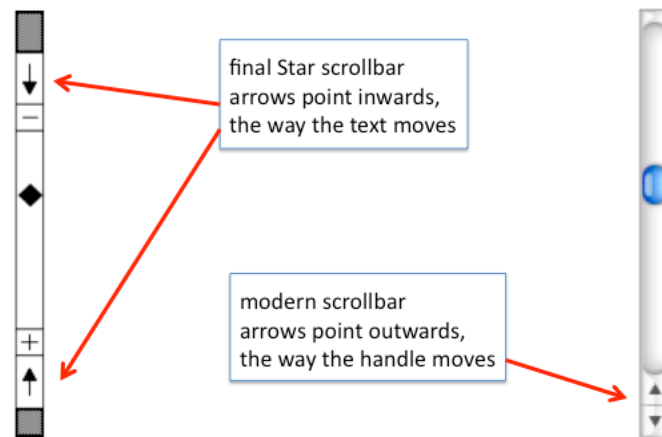


Figure 5. Xerox Star and modern (Mac OS X) scrollbars

So ideally, for good science, we would focus our evaluation where our justification is weakest, thus obtaining maximum information from our work and pushing forward the field. Of course, we should certainly be aware, while we probe these areas of greatest uncertainty, that our assumptions may be wrong, that the obvious may in fact turn out to be false; but we should not primarily make the obvious our focus.

There is of course a place, albeit largely absent in HCI research, for reproducing previous studies as a basis for further work, especially when the earlier work was promising but inconclusive. In mathematics you will go back and recheck the proofs of earlier theorems on which your work depends. If you do not do this and subsequently a flaw is found in the older proof then your own work fails with it. If such a flaw were found in, say, Nielsen and Landauer [21], would we have the means as a discipline to 'fail' all the succeeding work that relied on it?

There is a difference between reproducing previous studies for the purposes of verification and doing the obvious for the purposes of getting an 'easy hit'. The overarching aim should always be systematically to increase the knowledge of the field.

Sadly when advising students I have to tell them that there is a conflict between this recommendation for good science, and what is best to get published. The easiest way to get a publication is to choose something that you have a pretty strong argument for and then run some sort of experiment on it. With such an experiment you know what to expect and so you can frame a clear experimental hypothesis, and are very likely to get a result that will be statistically significant. However, this gives least new knowledge. In contrast, experiments focused on the weak points in the justification will have unknown answers, may yield inconclusive results and are least likely to have statistical power – while they have the potential to add knowledge to the discipline they are risky for the individual.

Note that this risk is the opposite of "nobody has done this, let's try it" experimentation. Instead it is the systematic exploration of gaps in knowledge set within a context that highlights valid possibilities.

As a discipline we should not find ourselves in a position where good science and publication are at odds. This is bad for new researchers entering the discipline and it is bad for the discipline itself. So when reviewing work we should seek

- (a) reasons why the issue/feature is as it is – that is, rational not just random ideas, *systematic* growth of the field;

and equally important

- (b) reasons why the issue/feature need not be as it is - that is, not obvious, *adding information* to the field.

Note that these two together ensure that collectively we systematically explore gaps in knowledge.

4.3 Roles

The discussion has moved to criteria for good work, but within HCI there are many genres of work, so we need different criteria of judgement depending on the genre. Again this can be a real problem during reviewing of papers. A couple of years ago as a meta-reviewer I had to explicitly say that I was entirely discounting one of the reviews, because the review was effectively criticising the genre of work within HCI, not assessing it with respect to criteria within the genre. Often it is not so clear and it is easy to let one's own general opinions about the most appropriate approach (experimental, ethnographic, formal) colour the judgement of a particular piece of work. Blandford [2] warns reviewers to judge research on its own merits, not "*would I have done it the same way*".

Now this is not to say that there is no place to critique and debate the validity of particular genres or approaches to work and assess their strengths and weaknesses when applied to particular problems. It is just that we should debate the validity of the genre as a whole within the discipline; and the validity of a given piece of work within its genre, so long as that genre is accepted as valid within the discipline and is applied within its understood bounds.

In the UK, the HCI community has noticed this is particularly problematic when it comes to reviews for projects and grants within HCI. It is very hard to get across-the-board support from reviewers to say a piece of work is of the best quality; someone will have something negative to say. If an HCI project is then viewed amongst those from different areas where the reviewing is more consistent (whether positive or negative), the best HCI projects will lose out compared to the best from those other areas. Now this is partly because the quality criteria within HCI are soft and less clear than in some areas, but partly because at least one of the reviewers may not like the general approach/genre.

There is certainly a need for discussion of the value of particular approaches and establishing new ones, but that should be a separate discussion. Furthermore, we need to think explicitly about these different approaches, techniques or genres and the criteria appropriate for each. This makes it easier to assign appropriate people to review work – and the recent CHI Conference subcommittees are an excellent move in this direction. Furthermore, if we are aware of these different criteria we can more easily say "personally I don't like this style of work, but within its genre it is strong".

Just as there are different genres of work, there are different roles that we may take within HCI research.

Imagine a physics paper that started off with some experiments at CERN, then performed group-theoretic analysis of superstring theory, and finally applied the results to the design of a vacuum cleaner. This is clearly risible. But HCI papers are often like this, and furthermore *expected to be*: a little bit of theory, build a toy system, run some

experiments, analyse the results, give implications for design⁷. Now this can sometimes be done well, so it is not that we should never have work like this, but surely it should be more common to have different aspects of this work performed by those that do them best, rather than expecting every paper to have a bit of everything?

HCI as an academic discipline (and maybe science) will develop most strongly if we can understand how different parts fit together and allow people and teams to focus where their core strengths lie.

I can think of three broad roles (although I am sure there are more):

- generating ideas and theories,
- developing systems and designs,
- performing empirical studies.

Table 1 and Figure 6 list some of the different criteria for each role; although in the process of delineating these criteria, empirical studies divide in two because the actual gathering of data may need different expertise from its analysis. This was certainly part of the origins of ethnography – report with as little interpretation as possible in order that someone else can interpret later.

Also, the parlous state of statistics in HCI [36] is no doubt in part due to the 'do it all' methods. In medicine there are special medical statisticians who are not medics themselves but do the statistics, because the medics themselves do not expect to be able to do this.

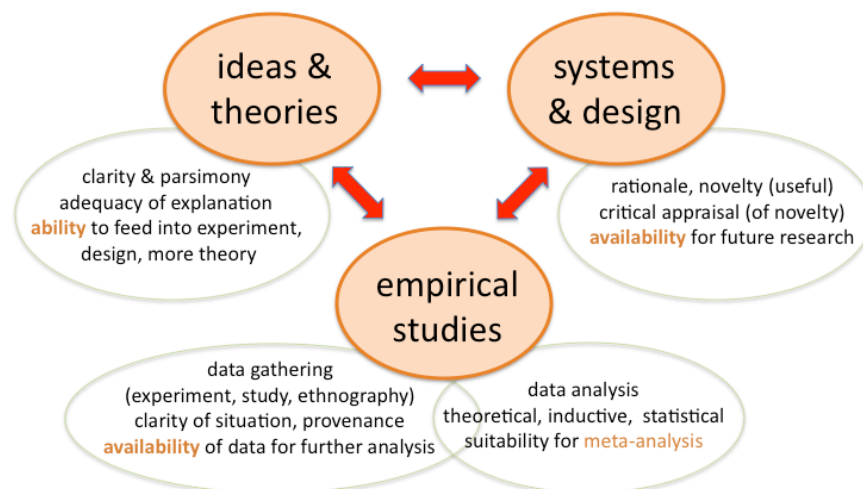


Figure 6. Roles in HCI research

⁷ See Dourish [37] for why 'implications for design' is not the ultimate measure of an effective paper.

Table 1. Roles and criteria

	some criteria
ideas & theories	clarity & parsimony adequacy of explanation ability to feed into experiment, design, more theory
systems and designs	rationale, novelty (useful) critical appraisal (of novelty) availability for future research
empirical studies – data gathering (experiment, study, ethnography)	clarity of situation, provenance availability of data for further analysis
empirical studies – data analysis (theoretical, inductive, statistical)	soundness, lack of bias suitability for meta-analysis

Within each role there are criteria that are more to do with internal coherence of the work, but, at least as important, there need to be criteria about making sure the work contributes to the bigger picture of the discipline as a whole.

- If I have a new theory or framework; is it expressed clearly enough so that someone else can apply it to their new design, or analysis of experimental results?
- If I have constructed a new system that embodies some idea; is it available so that other researchers can deploy it for long-term study, or use it in an experiment?
- If I have gathered some ethnographic data; have I obtained sufficient consents and described my data gathering techniques well enough so that the raw data can be made available for others to study in different ways?
- If I have performed some statistical analysis on an experiment; have I presented the results in a way that others can interpret and possibly perform meta-analysis?

The web was developed so that physicists could share data. We need to develop HCI so that we share data, systems and results equally easily, so that we can properly use each individual's expertise and skills to build a coherent discipline that is greater than any of us.

Long [38] emphasises the importance of a discipline having an accepted ontology, and in Long [8], effectively develops a framework providing just such a high-level ontology. This is effectively about common language in order to communicate clearly. Sutcliffe's [39] proposals for reuse of HCI knowledge, and those engaged in patterns research [40,41,42], instead seek to create common formats for sharing knowledge. In fact it is not essential that we all share a single language, for all our working and reporting, nor that we understand fully one another's methods. Within our sub-areas of HCI we can use our own esoteric languages, but the core outputs of our work need to be communicated clearly to enable others to build on them. It is not acceptable for an ethnographer to tell a technologist to go read Garfinkel, nor for me, with a PhD thesis originally about "Formal Methods and Interactive Systems", to tell an ethnographer to go read Gauss. Given the very broad nature of HCI, there may even be special roles for those who present the outcomes of one sub-area of HCI to others, a form of internal education within the discipline. Possibly we need suitable reward mechanisms such as special

high profile venues for such communication works, in the way ACM Computer Surveys served computing as a whole.

5 The changing face of HCI

Of course HCI is changing as computer technology changes and these changes will require yet deeper considerations of the way we interact together as an academic community and discipline. These changes also demand that we question more profoundly accepted methodological practice. Most readers will be aware of the rapid rate of change in recent years, but this section briefly reviews these changes, before looking at a specific case study of the way some of these changes forced methodological reflection.

The birth of HCI as a discipline was around the same time as the introduction of the desktop computer, and it became hard to consider interfaces that were not WIMP-based GUIs. This was to some extent a breaking free from the computer in the machine room, but it only got as far as the desktop, and there it stayed for nearly 20 years. Indeed, in 2001 Buxton wrote:

"In the early 1980s Xerox launched Star, the first commercial system with a Graphical User Interface (GUI) and the first to use the "desktop" metaphor to organise a user's interactions with the computer. Despite the perception of huge progress, from the perspective of design and usage models, there has been precious little progress in the intervening years. In the tradition of Rip van Winkle, a Macintosh user from 1984 who just awoke from a 17-year sleep would have no more trouble operating a "modern" PC than operating a modern car"
[43]

However, over the last 5 to 10 years (with plenty of preliminary research work before), we have seen the role of the computer change dramatically in society.

With mobile and ubiquitous computing and tangible interfaces, the computer physically escapes the desktop into the outside world. This is not just the subject of research for the future, but day-to-day reality for all of us. I often ask people "how many computers in your house" and most still say two or three or maybe (if they are very 'nerdy') more, but rarely do people remember their microwaves and HiFi, central heating and washing machines. Even our body load of computation is substantial. A few years ago I emptied my own pockets during a masters seminar and found four clear computers without anything particularly nerdy: (1) a mobile phone; (2) car keys, which include some form of coding processing for the remote locking; (3) a film camera, but with an LCD screen; and (4) the now ubiquitous chip in a credit card.




Figure 7. After 25 years chained to the desktop, the computer breaks free (images courtesy Matt Oppenheim)

The Internet has also seen the computer escape the desktop virtually (Fig. 7). While the growth of corporate networking led to the development of CSCW as an area, the Internet has been crucial in establishing collaboration that cuts across organisational boundaries and enters the home. It is also hard to believe that 10 years ago there was no Google. Equally it is easy to forget how un-ubiquitous universal Internet access is. In the 1998 business plan for aQtive, one of the dot.com companies I was involved with, we talked about designing products ready for the coming *PopuNet* [44] – the network for everyone, everywhere and everywhen. At the time this was just beginning to be a reality *when at the office desk*, but in the home a slow and expensive dial-up connection was the best one could hope for. The 'everywhen' was particularly critical – continuous not just continual, not just available (anytime), but always there; what is now termed "always on". For those with the iPhone in western cities this now appears to be the reality, but even within the UK, Europe or the US just move a little into rural areas, sail out to the islands, or climb amongst the mountains, and connectivity becomes more broken. Now move further afield to the developing world and we can see how easy it is to overestimate the universality of the Internet and correspondingly, albeit unwittingly, design to divide.

Amongst many things that have changed with the growth of the web is that much software that was a *product* has become a *service*. Although I still use an email client installed on my computer (because, whilst travelling a lot, I am not 'always on'), many only use web-based email services. Now we also have online word-processing, spreadsheets and more. If you buy an expensive hair styling kit, then you will continue to use it even if you find flaws, but if you visit a hairdresser and do not like the service or style, then the next time you go to a different one. Shrink-wrapped products allow you *one choice point* maybe every few years, but services allow near *continuous choice*. From a business perspective for shrink-wrapped software you can 'get away' with bad usability and poor user experience, so long as you have good marketing; the customers have already paid their money by the time they find out it is rubbish! In a service-based world, usability and user-experience become key to success.

As we can see, these technological changes lead to changes in the environment within which HCI works, not just because we have different hardware to play with, but because recent technological change has had a major impact at a commercial and social level. Sometimes technology is servant to social change, perhaps ignored or resisted. However, at various points technological changes have made a radical difference to social order. For example, the invention and adoption of the stirrup not only revolutionised mounted warfare, but was also a driver for the whole feudal system [45]. While wanting to avoid simplistic technological determinism, it is also clear that there are major societal, cultural and even cognitive changes that we need to recognise for

their impact on research and practice in HCI as well as for their broader political and ethical import.

One of the aspects that is obvious is the increasing focus within HCI on user experience. The physical movement of computers out of the office into the home and into our hands as well as the domestication of the web means that the old utilitarian 'efficiency and effectiveness' now have to pay second fiddle to 'satisfaction'. One of the amazing things about the numerous ethnographies of the home and of leisure (which are both methodologically harder than work ethnographies) is just how complex day-to-day life is [46]. The industrial revolution, Taylorism, and the continuing need to deal with staff turnover, has led to a largely controlled and 'normalised' workplace with personal differences minimised. Of course workplace studies constantly show how much fluid working depends on various adaptations, but always set against and located within a framework of order. In contrast the home has never needed the same levels of uniformity except those established externally by work and society.

The domestication of technology is nowhere more apparent than in Web 2.0 with its focus on user-contributed context and social networking [47]. Again it is interesting to look back to the dot.com days, less than a decade ago. In 1999, when working on a new product/service, vfridge, we articulated the idea of the *web sharer* [48]. At the time many were saying that there would be a shake-up in the web world with DIY home pages withering and all the traffic going to a tiny handful of sites (Yahoo!, AOL, Amazon) principally operating in a publication or broadcast manner, like the TV except with easier ways to buy things. Now this seems laughable, similar to the (misquoted) early predictions that five computers would be enough for the whole world⁸, but at the time was becoming accepted wisdom. In contrast we sought to design products for the 'web sharer': ordinary people sharing with one another.

Everyone may be a web sharer — not a publisher of formal public 'content', but personal or semi-private sharing of informal 'bits and pieces' with family, friends, local community and virtual communities such as fan clubs.

⁸ The "*five computers are enough*" quote is almost certainly a misquote. Some time ago the following question and answer were available on an IBM FAQ dated April 10, 2007, but this is no longer accessible in the IBM web site (although the FAQ text is quoted elsewhere). Whether or not Watson said it, it is certainly the case that no-one in the early days of the computer industry foresaw a future for computers outside very large corporations and government.

Q. Did Thomas Watson say in the 1950s that he foresaw a market potential for only five electronic computers?

A. We believe the statement that you attribute to Thomas Watson is a misunderstanding of remarks made at IBM's annual stockholders meeting on April 28, 1953. In referring specifically and only to the IBM 701 Electronic Data Processing Machine -- which had been introduced the year before as the company's first production computer designed for scientific calculations -- Thomas Watson, Jr., told stockholders that "IBM had developed a paper plan for such a machine and took this paper plan across the country to some 20 concerns that we thought could use such a machine. I would like to tell you that the machine rents for between \$12,000 and \$18,000 a month, so it was not the type of thing that could be sold from place to place. But, as a result of our trip, on which we expected to get orders for five machines, we came home with orders for 18."

This is not just a future for the cognoscenti, but for anyone who chats in the pub or wants to show granny in Scunthorpe the baby's first photos.

the web sharer vision [48]

There were probably others voicing similar ideas, but, like continental drift in the 1960s, this was against the prevailing wisdom of the time and so difficult to publish and hard now to trace. The crucial thing is that in 10 years what was completely counter-cultural has become passé. Given these rapid and substantial changes in the technical and social context of HCI, there is even greater need to re-examine methods and if necessary modify them or develop new methodology.

6 Case study: a single person study

Within this setting of changes in HCI we will focus on the PhD research of Fariza Razak on the use of single subjects in research and design. The purpose of this section is not to present the work in full, but rather to use it as a case study to illustrate the need for *methodological* thinking to address new kinds of issues. Because of this only sufficient details are presented for the purposes of illustrating the general issues; for more about Razak's work see her thesis "Single Person Study: Methodological Issues" [49]. Interestingly, while studying just one user initially seems to stand against all accepted HCI practice, in fact we shall see that single user studies of different forms are common in many disciplines.

6.1 Background

Razak started out interested in mobile user experience and in particular looking towards mobile learning. As a preliminary exercise she conducted a small study asking a handful of people about their use of mobile technology. As anyone who has done this sort of study knows, it is very hard to get beyond the banal – learning what you knew before you started. You ask questions and people tell you the answers you could have predicted, the difficulty is finding the questions to ask that are less obvious as questions and will lead to new knowledge – really growing the field.

In the initial study one of the participants stood out as unusual. She said she rarely used her mobile phone and yet other answers seemed to suggest the opposite, for example referring to use of time organisation features that others hardly mentioned. Clearly, her initial answer was about the 'normal' use of the phone as voice communication, but it was more to her than that.

Because this subject was unusual and different I suggested Razak spent some time investigating her in more detail. Little did either of us know at that time that this would become the key focus of Razak's PhD work.

Note that the subject was chosen because she was in one respect an extreme, an outlier, outside of the average. When analysing experimental results, outliers are often removed for statistical purposes, ignored as anomalous or extreme. Instead the outlier was chosen precisely because she was one. As an academic I always find extremes valuable. Partly because the abnormal, or extra-ordinary (strange how the words have different connotations) are just more interesting in themselves, but also because they cast light back onto the ordinary, showing us things that are often tacit and unnoticed. Djajadiningrat et al. [50] also describe how extreme characters helped expose aspects of use, especially 'undesirable' emotions and character traits, that more 'normal' personas and scenarios may hide.

Estrangement, the ability to see the world from a different perspective, is of great value in uncovering the hidden dimensions of the quotidian and I challenge students to have deliberate bad ideas [51]. Similarly, Merleau-Ponty [52] writes “*in order to see the world and grasp it as paradoxical, we must break with our familiar acceptance of it*” and schraefel [53] asks chemists to make cups of tea. Comedians are particularly good at this, seeing the oddness in the everyday, highlighting things that we recognise in ourselves that are slightly embarrassing or just strange when we look at them. Indeed I have often suggested that students look up humorous books about their domain of study as they may learn more from a comic’s eye than from many an academic study. The best ethnographers also seem to have this ability to see the normally overlooked details of a situation. Garfinkel [54] used ‘breaching experiments’ with his students, getting them to break normal social conventions in order to see that they exist; like a car engine, one is unaware of the parts until they fail.

This simple decision to study a single unusual subject started Razak down the route of pursuing this single person study as the central focus of her work and her thesis became not one about mobile experience or mobile learning, but a *methodological* account of the issues surrounding the use of a single person for research and design.

6.2 The first text

One of the early steps along this path was a diary study of Razak's subject. When we discussed the results of the study, the words of the very first entry leapt out of the page.

This first text read

*Dear God Don't need lots of frens! As long
as real ones stay with me, so bless them all,
especially the sweetest one reading this.*

and the subject’s comment (her emphasis):

this SMS MADE MY DAY!

Research on SMS behaviour often discusses its use for intimate communications “thinking of you”, “love you” [55,56]. However this message was something slightly, but significantly, different. The message here was in a way less personal; it was representative of a particular type of message: often small quotes of a devotional or otherwise encouraging nature, from friends, but not necessarily from her husband, or close family. John Rooksby (personal communication) described them as *messages that need no reply*. They are sent to encourage but not to establish communication in any interactive sense and certainly not to ‘communicate’ in an information sense – they are more gifts of thought.

Perhaps the closest thing in the physical world are the little cards or bookmarks that have poems, sayings or prayers written on them, often surrounded by flowers ... a world away from the design studio with its austere black-robed occupants.

This single text message and the reaction it caused fundamentally changed our view of the use of the mobile phone.

But can studying a single user in this way contribute to theoretical HCI research or practical interaction design?

6.3 Research from single-user studies

At first the idea of studying a single user runs counter to common academic sense. Surely we need to study many users to be able to stand any chance of generalising

results? However, it turns out that this is not so uncommon in other disciplines such as special education or studies of neurological deficit. For example Battro [57] studied the development of a child who had only a single brain hemisphere due to a congenital brain defect. Similarly, Damasio [58] in building up his understanding of the role of emotion in human reason, draws heavily on documentary evidence of Phineas Gage, who in 1848 suffered a traumatic brain injury leaving him with full intellectual capabilities but severe emotional deficit. In psychology there have been well respected uses of single subject studies, and even in HCI ethnographies are typically of a single situation (even if it includes several people) and, as we have seen, experiments often use a single application or piece of software (even if they have many subjects).

Furthermore, the study of a single user brings particular benefits. As often found in ethnographic studies, rich empirical data reveals new issues ... in this case the very first text! Furthermore, studying a single user in depth allows the researcher to build up a deep personal rapport with the subject and hence make sense of what would otherwise be irrelevant or meaningless aspects of the data.

In fact discovering *novelty* only needs *one example*, like a botanist discovering a new flower; a single specimen shows that the new species exists. Of course a different person at a different time in a different place would find a different flower. Studying a single person is not the way to find *all* the important issues, or establish *how common* a particular issue is, but the depth may be a good way to find new usage phenomena.

However, this still leaves us with the question: having found a new phenomenon, how common and critical is it? In other words, how do we generalise? In an empirical study if the sample of users is wide enough (not just psychology or CS students!), then we assume that if an issue is common in the sample it is common in the population. Of course, we can use new insights from any method, including those from studying a single user, to drive empirical work of this kind. However, in the case of the initial text message, extensive empirical studies were unnecessary for us to recognise that this was something that we would expect to see elsewhere, not for everyone (and maybe least for UK CS undergraduates, whom we might have studied in a larger scale survey), but at least for particular kinds of people and communities. That is we were able to *generalise by reasoning* based on the data we had seen, knowledge of other research work, our own personal experience and not least (albeit much undervalued in academia) common sense!

Generalisation through reasoning is again common in other areas, for example semiotics and mathematics, and is typically based in deduction or abduction rather than induction as used in reasoning from voluminous empirical data. However, possibly drawing on my mathematical roots, I would like to make a stronger claim:

generalisation never comes (solely) from data,

Instead, generalisation always comes through understanding. Even when we have copious data, the knowledge that we have chosen representative groups, the level of extrapolation we choose to make from the experimental tasks, or the belief in the methods are all matters of judgement. We generalise with our heads not our senses.

6.4 Designing for a single user

Akin to the research question of how we obtain knowledge from a single user is the practical one of whether we can use a single user in design. We all know that 'five users is enough' ... but one!

So as an experiment Razak attempted to design an application especially for that single user. Having got to know this individual intimately, what would be perfect for that single person? With a single user it is possible to spend sufficient time to collaboratively co-design in a way that is tuned for the specific lifestyle, abilities and personality of the user. Having done this one can then ask whether the application would work for others and maybe do more traditional user testing. Maybe this hyper-tuned application may form the start of a slightly more generalised application that is of more general appeal.

We do not expect such a perfectly tuned application to be liked by everyone, indeed often the opposite. In the case of Razak's subject the application periodically texted uplifting messages, and one test user clearly found some of the messages simply annoying. However, a surprising number of other users did find it engaging.

Again, like single-user studies in other disciplines, one does not have to go far to find areas where taking 100, 20 or even 5 users would seem like overkill ... indeed many designers find no users sufficient (although sometimes this is apparent).

In fact Nielsen and Landauer [21] calculate the “five users” figure based on a cost-benefit trade-off between the number of faults found with N users, the costs of performing the test with them, and the costs of a prototyping cycle. This calculation also took into account the fact that the number of new usability problems found with each additional user drops due to overlapping faults between users. If the costs of prototyping are high compared with the costs of usability tests, then it is worth doing more usability tests in each iteration cycle, if prototyping is cheaper or usability testing more expensive, then tighter cycles are optimal with fewer users tested per cycle. Nielsen and Landauer measured the *actual prototyping costs* in a number of projects compared with actual error rates and usability test costs and it was these empirical figures, *derived using 1993 technology and applications*, which gave rise to the now ubiquitous “five is enough”.

Of course prototyping costs are now substantially smaller than they were in 1993, and if the costs of prototyping are low enough, then the optimal point may not be five but even a single user per cycle. This is precisely the approach taken in Marty and Twidale's “extreme evaluation”, where short usability tests are carried out with a single person [59,60,61].

While extreme evaluation only evaluates with a single person *per prototyping iteration*, successive iterations will typically involve different people. Furthermore, the software being designed was not designed ‘for’ the single individual, just tested with one user. In contrast, Razak developed an application specifically optimised for just one user, although our expectation was that this would in fact lead to a concept that would have wider appeal. This form of single-user designing would not be good for all kinds of application of product, but is particularly useful when designing for *peak experience*.

6.4.1 Designing for Peak Experience

Imagine you have a group of children and want to give them lunch. In the UK you might well choose baked beans. Not the most exciting choice, but few children actively dislike baked beans; they are acceptable to everyone. However, give each of those children a euro (or maybe two) in a sweet shop ... they will all come away with a different chocolate bar, the chocolate bar that is 'OK' for everyone gets chosen by none. Or imagine choosing a menu for a wedding dinner ... maybe chicken with a bland sauce ... something nearly everyone will eat, but few people would choose for themselves in a restaurant.

Much of traditional HCI design is like baked beans – a word processor installed for the whole company, a mail program used by every student, good enough for everyone. However, increasing personal choice, especially for web-based services, makes design more like the chocolate bar; different people make different choices, but what matters is that the product chosen is not 'good enough' for all of them, but *best* for some.

This designing to be best for some, the chocolate bar rather than the baked beans, results in a product for *peak experience*. Figure 7 shows schematic profiles of three imaginary products. The horizontal axis represents different people/users and the vertical axis represents their level of satisfaction or experience. There is one 'good enough' product, which offers a consistent but mediocre experience. There are also two 'peak' products that offer high satisfaction for a few users and low satisfaction for others. Note with sufficient peak products the good-enough product will *never* be chosen.




Figure 7. Profiles of a 'baked beans' good-enough-for-all product, vs. two products offering peak of experience for some.

Traditional user-centred interface design may use user profiles, or personae chosen to be representative of a group as a whole, with a focus on the typical or the average – good for all. We typically move from identified user needs to interaction solutions with an emphasis on method and processes that ensure usability.

In contrast, designing for peak experience may need a stronger focus on the individual user, possibly extreme personae, focusing on the specific and eclectic – best for some. Often the move is from concept to use with an emphasis on novel ideas and inspiration.

Some years ago I was on a panel at ECCE with Jon Sykes from the group at Glasgow Caledonian studying games and emotion [62]. He was asked about the processes used by video games designers and many in the audience were shocked at the apparently *ad hoc* and non-user-centred way in which the designers have ideas, discuss them amongst themselves and only very late in the process submit them to user testing – but this is exactly what one would expect in order to design for peak experience; a good enough video game will be bought by no-one.

Similarly, many of the most successful Web 2.0 sites such as Facebook and del.icio.us started out being for the designers and their friends; tuned to a small group or even a single individual. We would normally castigate designers who design for themselves, but somehow, in spite of that, or maybe *because* of that, they are successful.

Interestingly, even the computer language of choice of many of these sites, PHP, was originally developed by one person for his own home page [63].

Now this is not to say that there is no role for traditional HCI practice; there are many products that do need to be used by everyone (e.g. bank web sites) and even in Web 2.0 web sites, such as YouTube, there are some aspects where traditional usability breaks down and experience dominates, but other parts, such as the uploading of videos, where standard usability is crucial [64].

However, where individual choice and user experience dominate, we need to look increasingly at peak experience. Mash-ups, widgets and open-data allow large numbers of applications designed for smaller groups of intended users; indeed one of the defining features of Web 2.0 [47] is this focus on the *long tail* of large numbers of web sites and web applications used by few people, as opposed to more traditional web applications aimed at mass use. We cannot expect that the vast army of mash-up builders will each employ a usability consultant; HCI for the long tail may need to consider how to build necessary aspects of usability into platforms or maybe even popularise HCI – the equivalent of 'house makeover' programmes on daytime TV.

7 Bringing it together

The single person study was introduced as a case study to show the importance of clear *methodological reflection*. At first it appears to flout the community conventions for effective HCI. However, by understanding methodology we were able to recognise similar methods in other disciplines and within HCI and also to see how it could be used effectively as part of research and design. Note that this is not to promote single-user studies above other techniques, but if it and other methods can be understood methodologically they can be applied where they are appropriate and give value depending on the context.

Note too that the adoption of the single user study was because we were addressing an issue at the changing boundaries of HCI. As a textbook writer I am always interested in what changes and what does not change between revisions. The things that have hardly changed in 15 years are likely to still be of value in another 15 years, but the things that changed in the last revision are likely to change again. We need to look particularly carefully at the new or changing things, so that we see the things of lasting value and are not simply swimming with fashion.

Just as taking the extreme user helped us to understand 'normal' use, so also as we look at new areas of technology, they help us understand afresh the old. Often the lens of unfamiliarity helps us explore the heart of things.

Others exploring the extremes of HCI have also needed to think seriously about methodology and how to adapt it to the circumstances of their work, for example, Button and Dourish [65], with technomethodology; McCarthy and Wright [11], in designing for experience and enchantment; and Gaver [66], in 'polyphonic assessment' of designs for everyday life.

The danger of establishing new methodology and potentially new vocabulary and theory is that we further fragment the genres and roles within HCI. How do we avoid a community of interest becoming a cabal?

This takes us back to the three challenges at the heart of this paper. We do not need to establish a common ontology or model for all of HCI, a single method that we all use; we

do not even need to understand the details of the language, methods and theories within all of our sub-communities. However, we do need to ensure that we understand the genres of work and the roles they play within a coherent discipline, we do need to ensure that the methods used within each role and genre are coherent, and above all we need to ensure that the results of each genre (not necessarily the full arguments and methods) are communicated in a way that is accessible to others in the wider community.

To be an academic discipline is about community, but not just any community, a community that establishes clear knowledge and together learns.

8 Acknowledgements

Many thanks to Aaron Quigley, Gavin Doherty and Liam Bannon for inviting me to share in the inaugural celebration of SIGCHI Ireland on which this paper was based; and to everyone who attended the presentation and chatted after and before that talk. Special thanks also to Fariza Razak whose work I have referred to extensively, to Matt Oppenheim (alias hardware monkey) for his wonderful cartoons of the computer breaking free, and to Fiona Dix for proofreading. Thanks also to the critical insights of the anonymous reviewers of this paper and to the special issue editors.

9 References

1. A. Dix, Human-Computer Interaction in the Early 21st Century: a Stable Discipline, a Nascent Science, and the Growth of the Long Tail, SIGCHI Ireland Inaugural Lecture, Trinity College Dublin, 2nd December 2008.
<http://www.hcibook.com/alan/talks/Dublin-2008/>
2. A. Blandford, HCI Research and Quality: Discussion Document. UCLIC/Equator two-day workshop on The Future of HCI in the UK: Research and Careers, (Loughborough University, 14-15th June, 2007).
<http://www.ucl.ac.uk/projects/future-uk-hci/>
3. A. Dix, Theoretical analysis and theory creation, Chapter 9 in: P. Cairns, A. Cox (Eds.), Research Methods for Human-Computer Interaction, Cambridge University Press, 2008.
4. Y. Rogers, New Theoretical Approaches For HCI, Annual Review of Information Science and Technology, 38 (2004).
5. B. Shackel, Ergonomics for a computer, Design 120 (1959) 36–39.
6. J. Long, J. Dowell, Conceptions of the discipline of HCI: craft, applied science, and engineering, in: A. Sutcliffe and L. Macaulay (Eds.), Proceedings of the Fifth Conference of the British Computer Society, Human-Computer interaction Specialist Group on People and Computers V (Univ. of Nottingham), Cambridge University Press, New York, 1989, pp. 9–32.
7. R. Harper, T. Rodden, Y. Rogers, A. Sellen, Being Human: Human-Computer Interaction in the Year 2020. Microsoft Research, Cambridge, 2008.
<http://research.microsoft.com/en-us/um/cambridge/projects/hci2020/>

8. J. Long, Specifying relations between research and the design of human-computer interactions, *Int. J. Human-Computer Studies*, 44 (1996) 875-920
9. D. Diaper, The discipline of human-computer interaction. *Interacting with Computers*, 1, 1 (1989) 3-5.
10. D. Diaper, C. Sanger, Tasks for and tasks in human-computer interaction. *Interacting with Computers*, 18, 1 (Jan. 2006) 117-138, DOI: <http://dx.doi.org/10.1016/j.intcom.2005.06.004>
11. J. McCarthy, P. Wright, *Technology as Experience*, MIT Press, Cambridge, 2004.
12. A. Dix, S. Gill, D. Ramduny-Ellis, J. Hare, Design and Physicality – Towards an understanding of physicality in design and use, in: *Designing for the 21st Century: Interdisciplinary Methods & Findings*. Gower, 2009.
13. R. Dawkins, Chapter 11, Memes: the new replicators. in: *The Selfish Gene*, Oxford University Press, London, 1976
14. H. Clark, S. Brennan, Grounding in communication, in: L. Resnick, J. Levine, S. Teasley (Eds.), *Perspectives on Socially Shared Cognition*, American Psychological Association, Washington, 1991, pp. 127-149.
15. H. Clark, *Using Language*. Cambridge University Press, Cambridge, 1996.
16. S. Star, The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving, in: L. Gasser, M. Huhns (Eds.), *Distributed Artificial Intelligence, Volume II*, Morgan Kaufmann, SF Mateo, 1989, pp. 37-54.
17. A. Dix, J. Finlay, G. Abowd, R. Beale, Interaction design basics, Chapter 5 in: *Human-Computer Interaction*, third edition. Prentice Hall, 2004
18. S. Brand, *How Buildings Learn: What Happens After They're Built*, Viking Press, 1994.
19. J. Nielsen, Why You Only Need to Test With 5 Users, Alertbox, March 19, 2000. <http://www.useit.com/alertbox/20000319.html>
20. G. Miller, The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information, *The Psychological Review*, 63 (1956) pp. 81-97. <http://www.musanim.com/miller1956/>
21. J. Nielsen, T. Landauer, A mathematical model of the finding of usability problems, in: *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (Amsterdam, The Netherlands, April 24 - 29, 1993). ACM Press, New York, 1993, pp. 206-213. <http://doi.acm.org/10.1145/169059.169166>
22. K. Larson, M. Czerwinski, M. (). Web Page Design: Implications of Memory, Structure and Scent for Information Retrieval, in: *Proceedings of CHI98, Human Factors in Computing Systems*, ACM Press, 1998, pp. 25-32.
23. A. Dix, Upside down As and algorithms – computational formalisms and theory, Chapter 14 in: J. Carroll (Ed.), *HCI Models, Theories, and Frameworks: Toward a Multidisciplinary Science*, Morgan Kaufmann, San Francisco, 2003, pp. 381-429. <http://www.hcibook.com/alan/papers/theory-formal-2003>

24. J-F. Rouet, C. Ros, G. Jégou, S. Metta, Locating relevant categories in web menus: effects of menu structure, aging and task complexity, in: D. Harris, V. Duffy, M. Smith, C. Stephandis (Eds.), *Human-Centred Computing: Cognitive Social and Ergonomic Aspects*. Vol. 3 of *Proc. of HCI Intl* Laurence Earlbaum, New Jersey, 2003, pp. 547–551.
25. B. Eisenberg, Debunking Miller's Magic 7. ClickZ, Oct 29, 2004.
<http://www.clickz.com/3427631>
26. A. Dix, Controversy and Provocation. In *Proceedings of HCIE2004, The 7th Educators Workshop: Effective Teaching and Training in HCI*, (Preston, UK 1st and 2nd April 2004). ISBN 0-9541927-5-3
<http://www.hcibook.com/alan/papers/HCIE2004/>
27. A. Dix, European HCI Theory – a uniquely disparate perspective, *European HCI Research special area CHI 2004*, Vienna, Austria, 24-29 April 2004.
<http://www.hcibook.com/alan/papers/chi2004-euro-theory/>
28. A. Dix, Teaching Innovation, in: *Excellence in Education and Training convention*, Singapore Polytechnic, 17th May 2002.
<http://www.hcibook.com/alan/talks/singapore2002/>
29. G. Ellis, A. Dix, An explorative analysis of user evaluation studies in information visualisation, in: *Proceedings of the 2006 Conference on Beyond Time and Errors: Novel Evaluation Methods For information Visualization* (Venice, Italy, May 23 - 23, 2006), BELIV '06, ACM Press, New York, 2006, pp. 1–7.
<http://www.hcibook.com/alan/papers/beliv06-evaluation/>
30. D. Furniss, A. Blandford, P. Curzon, Usability evaluation methods in practice: understanding the context in which they are embedded, in: *Proceedings of the 14th European Conference on Cognitive Ergonomics: Invent! Explore!* (London, August 28–31, 2007). ECCE '07, vol. 250. ACM Press, New York, 2007, pp. 253-256. DOI: <http://doi.acm.org/10.1145/1362550.1362602>
31. A. Blandford, J. Hyde, T. Green, I. Connell, Scoping analytical usability evaluation methods: a case study, *Human–Computer Interaction*, 23 (2008) 278–327. DOI: 10.1080/07370020802278254
32. K. Popper, *The Logic of Scientific Discovery*, Basic Books, New York, 1959.
33. J. Johnson, T. Roberts, W. Verplank, D. Smith, C. Irby, M. Beard, K. Mackey, The Xerox Star: A Retrospective, *Computer*, 22, 9 (Sep. 1989) 11–26, 28–29.
<http://dx.doi.org/10.1109/2.35211>
34. A. Dix, Hands Across the Screen – why scrollbars are on the right and other stories, *Interfaces*, 37 (Spring 1998) 19–22,
<http://www.hcibook.com/alan/papers/scrollbar/>
35. A. Dix, Sinister Scrollbar in the Xerox Star Xplained. *Interfaces*, 38 (Summer 1998) 11, (short update to the above article),
<http://www.hcibook.com/alan/papers/scrollbar/scrollbar2.html>
36. P. Cairns, HCI... not as it should be: inferential statistics in HCI research, in: L. Ball, M. Sasse, C. Sas, et al (Eds.), *Proc. of HCI 2007*, Vol 1 BCS, 2007, pp. 195–201
37. P. Dourish, Implications for design, in: R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, G. Olson (Eds.), *Proceedings of CHI'06, the SIGCHI Conference on Human*

- Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006), ACM Press, New York, 2006, pp. 541-550. DOI:10.1145/1124772.1124855
38. J. Long, Theory in Human-Computer Interaction? in: IEE colloquium on Theory in Human-Computer Interaction (HCI), IEE, London, (Digest No.192), 17 Dec 1991, pp. 2/1-2/6.
<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=241136&isnumber=6182>
 39. A. Sutcliffe, On the effective use and reuse of HCI knowledge, ACM Trans. Comput.-Hum. Interact. 7, 2 (Jun. 2000), 197-221. DOI=
<http://doi.acm.org/10.1145/353485.353488>
 40. J. Tidwell, Common Ground: a pattern language for human-computer interface design, dated 1999, accessed May 2009.
http://www.mit.edu/~jtidwell/interaction_patterns.html
 41. J. Tidwell, Designing Interfaces: Patterns for Effective Interaction Design, O'Reilly, 2005.
 42. M. van Welie, Welie.com: Patterns in Interaction Design, dated 2008, accessed May 2009. <http://www.welie.com/>
 43. W. Buxton, Less is More (More or Less), in: P. Denning (Ed.), The Invisible Future: The seamless integration of technology in everyday life, McGraw Hill, New York, 2001, pp. 145-179.
 44. A. Dix, PopuNET: pervasive, permanent access to the Internet, eBulletin, aQtive ltd., July 1998. <http://www.hiraeth.com/alan/ebulletin/PopuNET/PopuNET.html>
 45. L. White, Medieval Technology and Social Change, Oxford University Press, 1968
 46. M. De Certeau, The Practice of Everyday Life, trans. S. Rendall, University of California Press, Berkeley, 1984. original French: L'invention du Quotidien. Vol. 1, Arts de Faire', 1980.
 47. T. O'Reilly, What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software, O'Reilly Media, 30th Sept. 2005, (accessed 23/06/2007)
<http://www.oreillynnet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>
 48. A. Dix, The Web Sharer Vision, eBulletin, aQtive ltd., November 1999.
<http://www.hiraeth.com/alan/ebulletin/websharer/>
 49. F. Razak, Single Person Study: Methodological Issues, PhD Thesis. Computing Department, Lancaster University, UK, February 2008.
<http://www.hcibook.net/people/Fariza/>
 50. J. Djajadiningrat, W. Gaver, J. Frens, Interaction relabelling and extreme characters: Methods for exploring aesthetic interactions, in: D. Boyarski, W. Kellogg, (Eds.), Proceedings of DIS2000 Designing Interactive Systems: Processes, Practices Methods and Techniques (New York 17-19 August 2000), ACM Press, New York, 2000, pp. 66-71.
 51. A. Dix, T. Ormerod, M. Twidale, C. Sas, P. Gomes da Silva, L. McKnight, Why bad ideas are a good idea. in: Proceedings of HCIEd.2006-1 Inventivity, Ballina/Killaloe, Ireland, 23-24 March 2006.
<http://www.hcibook.com/alan/papers/HCIEd2006-badideas/>

52. M. Merleau-Ponty, *M. Phénoménologie de la Perception*, Gallimard, Paris, 1945. (quote in text from: M. Merleau-Ponty, *Phenomenology of Perception*, K. Paul (trans.), Routledge, 1962, 2002)
53. m. c. schraefel, A. Dix, Within bounds and between domains: reflecting on making tea within the context of design elicitation methods, *International Journal of Human-Computer Studies*, 67, 4 (April 2009) 313–323.
54. H. Garfinkel, *Studies in Ethnomethodology*. Prentice- Hall, Englewood Cliffs, 1967.
55. L. Gamberini, A. Spagnolli, P. Pretto, Temporal structure of SMS-mediated conversation. in: *Time Design Workshop, CHI2004*, Wien, April 25, 2004.
56. A. Spagnolli, L. Gamberini, Interacting via SMS: practises of social closeness and reciprocation, *British Journal of Social Psychology* 22 (2007) 343–364.
57. A. Battro, *Half a Brain is Enough: The Story of Nico*. Cambridge Studies in Cognitive and Perceptual Development (No. 5), Cambridge University Press, 2001.
58. A. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*, Putnam Publishing, 1994 (paperback: Penguin, 2005)
59. P. Marty, M. Twidale, Lost in gallery space: A conceptual framework for analyzing the usability flaws of museum Web sites, *First Monday*, 9, 9 (2004).
60. P. Marty, M. Twidale, *Extreme Discount Usability Engineering*, Technical Report ISRN UIUCLIS--2005/1+CSCW.
61. P. Marty, M. Twidale, *Usability@90mph: Presenting and Evaluating a New, High-Speed Method for Demonstrating User Testing in Front of an Audience*, *First Monday*, 10, 7 (July 2005).
62. J. Sykes, presentation at panel on “Funology: A Science of Enjoyable Technology?”, M. Blyth (chair), ECCE-12, (St Williams College, Yor, 12-15 Sept. 2004)
63. PHP Manual: A History of PHP, accessed 1st May 2009.
<http://gtk.php.net/manual1/en/html/intro.whatis.php.history.html>
64. P. Silva, A. Dix, Usability – Not as we know it! in: *Proceedings of BCS HCI 2007, People and Computers XXI*, BCS eWiC. 2007.
<http://www.hcibook.com/alan/papers/HCI2007-YouTube/>
65. G. Button, P. Dourish, Technomethodology: paradoxes and possibilities, in: M. Tauber (Ed.). *Proceedings of CHI '96, The SIGCHI Conference on Human Factors in Computing Systems: Common Ground* (Vancouver, British Columbia, Canada, April 13 - 18, 1996). ACM, New York, NY, 19-26. DOI=
<http://doi.acm.org/10.1145/238386.238394>
66. W. Gaver, Cultural commentators: Non-native interpretations as resources for polyphonic assessment, *Int. J. Human-Computer Studies* 65 (2007) 292–305