A Preliminary Analysis of the Products of HCI Research, Using *Pro Forma* Abstracts

William Newman

Rank Xerox Research Centre
61 Regent Street
Cambridge CB2 1AB, UK
E-mail: newman@europarc.xerox.com

ABSTRACT

A classification scheme for the products of engineering research is described, involving three principal categories of product: improved modelling techniques, solutions and tools. These categories can be linked to the contributions they make to engineering design. A set of *pro forma abstracts* are proposed as a reliable means of identifying the three categories. A preliminary sample of published engineering papers indicates that normally at least 90 percent of papers fall into these three categories. For recent CHI and InterCHI conferences, however, only about 30 percent can be thus categorized. The remainder appear mostly to describe *radical solutions* (solutions not derived from incremental improvements to solutions to the same problem), and *experience and/or heuristics* gained mostly from studies of radical solutions. Some comments are made about the reasons for these departures from normal engineering research practice.

KEYWORDS: Human-computer interaction, research methods, research products, system design, abstracts, radical solutions.

INTRODUCTION

How will we judge CHI '94 in the years to come? By what metrics or criteria will we assess the value of this conference, or indeed of any conference relating to our fields of work?

Reprinted from *Proceedings of CHI '94 Human Factors in Computing Systems* (April 24-28, Boston, MA) ACM/SIGCHI, N.Y., pp. 278-284.

These questions of course elicit different answers from those who have attended the conference in question and those who have not. Attending a conference offers the unique benefits of meeting old acquaintances and making new ones, hearing topical issues addressed by keynote speakers and panels of experts, participating in workshops, and joining in the discussions of presented papers. It also offers a chance to learn about recent research from the papers themselves. None of these benefits, except the last, is available directly to non-attendees. For them the primary source of benefit is the set of papers published in the proceedings, and the measure of a good conference – perhaps worth attending next year? – is likely to lie in the work that the individual papers present.

I have become interested in understanding how published HCI research contributes to strengthening the practice of computer systems development. This is of course only a small part of the much larger issue of how all research, published or not, makes contributions. However, published research represents the primary basis of dissemination to the systems community as a whole, and therefore plays a vital role in strengthening practice. It is this role of published research, together with its accessibility, that has motivated me to conduct this study. As I shall point out, mine has been a very preliminary study, but it perhaps makes the case for further studies of a more extensive and rigorous nature.

This paper analyses HCI research in terms of the different kinds of *product* that result. In this respect it takes a different line from previous studies such as Jarvinen's which are more concerned with the researcher's method [14]. In outline, the paper starts by identifying a set of three principal categories of research product, corresponding to enhancements in modelling techniques, in solutions and in design tools. It verifies that these three categories do indeed cover the vast majority of research in other fields of engineering. It then considers the products of HCI research, pointing out that these are more difficult to classify. A more rigorous method of classification, based on *pro forma abstracts*, is developed, and this is applied to research publications in CHI and InterCHI proceedings. A relatively small proportion of these fall into the three principal categories; two further categories (radical solutions and experience/heuristics) are needed in order to cover the remainder, and the reasons for this are discussed in the final section.

ENGINEERING RESEARCH AND DESIGN PRACTICE

HCI research takes place within an engineering context. Its primary value lies in its contributions to the practice of interactive computer systems development. Leading HCI researchers have confirmed the link with engineering, justifying their work in terms of having "simplified the theory into practical engineering models, which are the tools for designers to apply the theory" [2]. I have attempted to place this study in the general context of engineering research, looking specifically at how this research can contribute to engineering *design*.

The engineering design context is especially vital to HCI research, for it defines not only a community of designers as clients for the research, but also a whole tradition of practice within which the results of research are used. We can find numerous descriptions of this practice in the literature [10, 11]. Taken together, they offer roughly the following account of the engineering design process:

1. Recognising the need for an artifice, and thus identifying a problem in computer systems design whose solution will meet this need.
2. Thinking of a number of alternative solutions to the problem, and making a selection.
3. Working out the details to see if the solution is practicable, calling on relevant technologies and engineering sciences.
4. Identifying aspects of the problem that the solution cannot address without additional knowledge, i.e., where improvements to the chosen solution are needed.
5. Making estimates and predictions of performance to see if the solution meets the specification.
6. Building a prototype and testing it thoroughly.
7. Making further modifications to the design, and possibly building and testing further prototypes.

The value of this account, indeed of almost any account of engineering design, is that it shows us where research can contribute. It helps us to identify three main forms of contribution that research can make:

EM *Enhanced analytical modelling techniques*, based on relevant theory, that can be used to tell whether the design is practicable or to make performance predictions;
ES *Enhanced solutions* that overcome otherwise insoluble aspects of problems, or that are easier to analyse with existing modelling techniques;
ET *Enhanced tools and methods* for applying analytical models and for building functional models or prototypes.

We can find confirmation of this classification in texts on engineering design, such as [12]. However, the real proof of its validity lies in the research literature. The results of a preliminary survey of five bodies of published engineering research, summarised in Table 1, suggest that at least 90 percent of papers normally fall into one or other of the three categories.

In summary, existing definitions and descriptions of engineering design enable us to make two observations. First they help us conclude that engineering research generally makes contributions to practice of three kinds: modelling techniques, enhanced solutions, and tools and methods. Second, they help us to see how each of these has a clearly identified role in support of the engineering design process. With these two points in mind, we can turn to the main question raised at the outset: what kinds of contribution does HCI offer to the practising engineer? The next two sections explain how we can overcome some initial difficulties in classifying HCI research. They describe the use of *pro forma abstracts* to capture the product of the research and set it in a context of the method by which it is carried out.

	EM	*ES*	*ET*	*Other*	*Tot.*	%*E*..
Electronics	51	28	12	1	92	99
Thermodynamics	82	2	--	--	84	100
Nuclear technology	79	19	20	10	128	92
Aeronautics	99	5	23	7	134	95
Soil mechanics	102	2	4	4	112	96
Total	417	56	59	22	554	96

Table 1. Numbers of published papers in each of three categories: Enhanced Model (*EM*), Enhanced Solution (*ES*) and Enhanced Tool (*ET*). The "%*E*.." column shows the percentage falling in these three categories. Sources: *Intl. J. Electronics* **70**, **71**; *Intl. J. Heat and Mass Transfer* **35**; *Nuclear Technology* **97**; *J. Aircraft* **26**; *12th Intl. Conf. on Soil Mech. and Foundations Eng*.

CONTRIBUTIONS FROM HCI RESEARCH

I have found HCI's research products difficult to classify in comparison with other branches of engineering. Although some published papers seem to be offering results that would fall into one or other of the three categories above, other papers clearly are not, and the dividing lines are hard to establish. For these reasons I have concluded that it is not possible to analyse the HCI literature by looking simply at the titles of papers, for example, or at the abstracts provided by authors. Such an approach is bound to be heavily influenced by the investigator's subjective opinion.

I have chosen to solve the problem of classifying ambiguous cases by trying to match them with the research methods employed in each of the three principal categories:

- *Enhanced modelling techniques* are generated by applying existing techniques to designs, and making comparisons with the results of empirical tests on working prototypes or products; areas of discrepancy between predicted and actual results are noted. The researcher then develops a new model that offers a more accurate prediction, and confirms its accuracy by means of fresh empirical tests.
- *Enhanced solutions* arise from measuring limitations in the ability of existing solutions to address certain problems. These motivate the researcher to devise an enhanced design that overcomes the limitations, perhaps making use of modelling techniques to predict the outcome. The new solution is prototyped and tested to demonstrate a successful outcome.
- *Enhanced tools and methods* arise from observations that modelling techniques or solutions need supporting tools or methods in order to be applied efficiently and reliably. A tool or method is devised, and is applied in a design context so as to confirm that it provides effective support for the modelling technique or for the use of the solution.

Each of the three methods, whilst following a similar pattern, stands out from the others in a number of ways. The overall pattern is one of initial analysis, motivating the devising of an enhancement, ultimately confirmed in tests. However, not only does each method devise a different kind of enhancement, but each initial analysis has a different focus, and the final confirmation has a different form too. In effect the conduct of the research stamps it with a methodological "signature" that sets it clearly apart from the other categories. It is this signature that *pro forma* abstracts attempt to capture, thus making each of the research categories more easily distinguishable.

Pro Forma **Abstracts**

Pro forma abstracts are templates, written in the style of normal abstracts, into which the results of research can be "slotted" according to the category of method followed and research product generated. There are, not surprisingly, three principal *pro formas* corresponding to the three types of product identified above:

EM Enhanced Model:

Existing <model-type> models are deficient in dealing with <properties> of <solution strategy>. An enhanced <model-type> is described, capable of providing more accurate analyses / predictions of <properties> in <solution strategy> designs. The model has been tested by comparing analyses / predictions with empirically measured values of <properties>.

ES Enhanced Solution:

Studies of existing <artefact-type> have shown deficiencies in <property>. An enhanced design for an <artefact-type> is described, based on <solution strategy>. In comparison with existing solutions, it offers enhanced levels of <property>, according to analyses based on <model-type>. These improvements have been confirmed / demonstrated in tests of a working <artefact-type> based on the design.

ET Enhanced Tool:

The effectiveness of <model-type> / <solution strategy> in supporting the design of <artefact-type> has been demonstrated. An enhanced tool / method is described for the design of <artefact-type> based on <model-type> / <solution strategy>. Examples are provided confirming the effectiveness of its support for <model-type> / <solution strategy> in design.

Examples of the use of these three *pro formas* to capture the products of HCI research are shown in Figure 1.

Enhanced Model:

Existing GOMS models are deficient in dealing with the speed of use of workstation applications involving dynamic visual information and multi-party conversation. An enhanced GOMS model is described, capable of providing more accurate predictions of speed of use in such workstation application designs. The model has been tested by comparing predictions with empirically measured values of speed of use (John90).

Enhanced Solution:

Studies of existing automatic document layout schemes have shown deficiencies in ease of learning and range of information handled. An enhanced design for a layout system is described, based on morphological analysis to extract logical structure. In comparison with existing solutions, it offers enhanced levels of accuracy in determining logical structure. These improvements have been demonstrated in tests of a working layout system based on the design (Iwai89).

Enhanced Tool:

The effectiveness of walkthroughs in supporting the design of interactive systems has been demonstrated. An enhanced tool is described for the design of interactive systems based on the use of video recording equipment and informal, interactive evaluation sessions. Examples are provided confirming the effectiveness of its support for walkthroughs in design (Rowley92).

Figure 1. Examples of the three principal forms of *pro forma* abstract, applied to published accounts of HCI research. Sources: John, *CHI '90 Proc.*, pp. 107-114; Iwai *et al.*, *CHI '89 Proc.*, pp. 369-374; Rowley and Rhoades, *CHI '92 Proc.*, pp. 389-395.

The examples in Figure 1 adhere quite closely to the *pro formas,* but I have not found this strictly necessary for the purposes of classification. It is possible to omit the motivational or confirmatory clauses altogether if the paper's contents warrant this.

HCI classification: initial results

Table 2 shows the initial results of categorizing the entire body of papers from the five sets of CHI and InterCHI proceedings published in the period 1989 to 1993, using the three principal *pro formas*. The results are significantly different from any gained from analysing research products from other fields. The most striking difference is the very high proportion of products falling into the "Other" category: out of 282 papers, 208 were "Others", representing nearly 74 per cent. In other words, roughly three quarters of the research presented at CHI and InterCHI conferences basically appears not to fit into any of the three principal categories of research. It simply resists classification according to normal engineering research practice.

We may be tempted to draw some immediate conclusions about the value of the research in question. However, a more important task is to learn more about the nature of the "unclassifiable" three quarters of the papers – what kinds of products did they offer, and what contributions might they have been intended to make to design? In the next section I will summarize the results of an analysis that led me to conclude that these papers were primarily concerned with presenting and analysing *radical solutions* that could not easily be related to previously published research.

Radical and normal solutions

Several accounts of engineering design have drawn attention to the special role of *radical solutions* in punctuating the normal steady progress of design enhancement. Constant draws a distinction between *radical technology* and *normal technology,* defining normal technology as "what technological communities usually do" [3]. He considers it generally requisite to technological progress, constituting the bulk of all technical activity.

Writing more recently, Vincenti contrasts normal with radical design:

> In radical design, how the device should be arranged or even how it works is largely unknown. The designer has never seen such a device before and has no presumption of success. The problem is to design something that will function well enough to warrant further development [14].

I have found it helpful to apply the radical-normal distinction to HCI research. Applied in the strict sense, very few of the systems described would count as "radical technology" or "radical designs", but a great many of them are distinctly novel in the way they apply technology to solving problems. They are truly *radical solutions,* of the kind to which Rogers is referring when he says, "a new concept may be an original application of an existing device or principle which necessitates a good deal of adaptation" [11]. In other words, often it is the *application* of interactive technology that is radical, not the technology itself. Among the many radical solutions described in the pages of CHI and InterCHI proceedings are the gaze-directed, self-disclosing display of Starker and Bolt, the shared awareness tool of Borning and Travers, and the paper user interface of Johnson et al. [1, 5, 13].

How can we reliably recognise radical solutions? Thomas Kuhn, who first drew attention to the distinction between revolutionary and normal science, suggested that normal science follows established *paradigms* [6]. At the time these paradigms are first proposed, they represent the kinds of revolutionary ideas that we see in radical design solutions. In an engineering domain they are characterized, as Constant points out, by a tendency to introduce perturbations and side-effects which must gradually be ironed out, and which make it difficult to measure improvements in any one dimension [3]. In HCI research, radical solutions stand out as deliberate attempts to introduce new paradigms in order to solve particular problems; they are characterized by the introduction of side-effects and by a need to reorganize the receiving environment. They are justified, nevertheless, by the possibility of overcoming the limitations of existing normal solutions.

These characteristics of radical solutions make them fairly easy to distinguish in the research literature. Many of them are explicitly described as "new paradigms," as "new conceptual frameworks for design," and so forth. They are also quite easy to identify by the simple technique of trying to force them into the Enhanced Solution *pro forma*. From the experience of trying, I was led to the following additional *pro forma* for describing radical solutions:

RS Radical Solution:

A radical solution to the problem of <problem definition> is described, based on <solution strategy>. In comparison with <existing normal solutions> it offers <advantages>, which have been demonstrated in prelim-

	EM	*ES*	*ET*	*Other*	*Total*	*%E..*
CHI '89	3	3	13	35	54	35.1
CHI '90	5	2	4	36	47	23.4
CHI '91	3	1	7	41	52	21.1
CHI '92	9	1	8	49	67	26.9
InterCHI '93	4	1	10	47	62	24.2
Totals	24	8	42	208	282	26.2
percent	8.5	2.8	14.9	73.8	100.0	

Table 2. Initial analysis of HCI papers, showing distribution of CHI and InterCHI papers among the three principal categories of product. As before, the "%*E..*" column shows the total percentage falling in these three principal categories.

inary tests, but it leaves a number of side-effects to be addressed including <list of side-effects>. Strategies are suggested for addressing these side-effects.

Of the 282 papers classified, 90 described solutions of some form, and of these only 8 could be clearly identified as enhanced solutions, while 70 clearly fitted the radical-solution *pro forma*. The remaining 12 fell in between, mostly because they could not offer any analytical or empirical evidence to justify their claims. However, since all of these borderline cases described attempts to improve existing paradigms, I chose to include them in the Enhanced Solution category. Figure 2 gives examples of both Radical Solution and "borderline" Enhanced Solution *pro forma* abstracts.

Design Experience and Heuristics

One remaining issue in HCI research is how radical solutions make the transition to normal practice, and what kinds of research aid this transition. What happens next after a radical solution has been invented? Obviously, many such solutions fall by the wayside, but some of them are eventually accepted into normal practice, and are then taken through stages of progressive enhancement. But until they gain acceptance, to what kinds of research do they owe their progress, and can we find the products of this research reported in the literature?

As we have seen, radical solutions bring about change along many dimensions, often in unpredictable ways. In the early days of their use, therefore, it is hard to study any single aspect of their behaviour. When studies are carried out, they tend to arrive at a collection of qualitative observations rather than a single quantitative one. We therefore see many examples of broad-ranging studies of new solutions which do not attempt to focus on any one issue, but present the findings as a whole. As the studies become more focused, they may seek to encapsulate observations into guidelines or heuristics for the benefit of designers. Ultimately we should expect them to reach a point where improved modelling techniques or enhanced solutions (borderline or otherwise) start to emerge.

Studies leading to publication of experience or design heuristics can be captured in abstracts based on the following *pro forma*, examples of which are shown in Figure 3:

XH Experience and/or Heuristic
Studies reported here of <application> supported by <supporting technology> generate a number of findings concerning <issues>, including <list-of-findings>. They indicate that <requirement> is / is not met by <design-heuristic>.

Summary of results of survey

After the analysis of the "Other" papers from the initial survey, which identified those covering Radical Solutions or Experience/Heuristics, and reclassified the "borderline" Enhanced Solutions, the results were as shown in Table 3. The proportion of papers falling outside the three traditional categories is still high, at around 70 percent. The most common are reports of experience and heuristics, which account for 113 of the 282 papers covered by the study, or over 40 percent. Radical solutions, at 70 out of 282 or roughly 25 percent, are the next most frequent. Of

---

(a) A radical solution to the problem of information navigation is described, based on analysis of eye movement to infer which item holds most interest for the user and therefore merits exploration in more detail (Starker90).

(b) A radical solution to the problem of rapid, reliable text entry in pen-based systems is described, based on an alphabet of single-stroke characters or *unistrokes*. In comparison to ordinary hand printing it offers increased speed, fewer recognition errors and "eyes-free" entry, which have been demonstrated in preliminary tests, but it leaves a number of side-effects to be addressed including difficulties in meeting requirements for rapid learning and large alphabets. Strategies for overcoming these difficulties are suggested (Goldberg93).

(c) Studies of existing text-cursor positioning techniques have shown deficiencies in speed and error rates. Enhanced designs for cursor positioning and display are described, based on a special shift key for moving to a text pattern and separate display of insertion and deletion points. In comparison with existing solutions, they offer enhanced speed of cursor movement and fewer errors (Raskin89).

Figure 2. Examples of Radical Solution *pro formas* (a and b) and "borderline" Enhanced Solutions (c). Sources: Starker and Bolt, *CHI '90 Proc.*, pp. 3-9; Goldberg and Richardson, *InterCHI '93 Proc.*, pp. 80-87; Raskin, *CHI '89 Proc*., pp. 167-170.

Studies reported here of collaboration supported by a video environment generate a number of findings concerning shared work support, privacy, awareness and blurring of boundaries, including the need to support the full range of shared work and to meet users' desire for both privacy and unobtrusive awareness (Gaver92).

Studies reported here of searching for documents supported by graphical icons generate a number of findings concerning search times. They indicate that speed of visual searching is best met by using simple icon designs (Byrne93).

Figure 3. Examples of Experience/Heuristics *pro formas*. Sources: Gaver *et al*., *CHI '90 Proc.*, pp. 27-35; Byrne, *InterCHI '93 Proc.*, pp. 449-453.

the three traditional categories, Enhanced Tools and Methods are most frequently reported.

DISCUSSION

A number of questions are raised by this study. Why, for example, do HCI researchers invest so much effort in developing and studying radical solutions? Is it beneficial to the systems development industry that they should do so? I conclude this paper with brief discussions of these two questions, but turn first to two other matters: the conduct of the study itself, and some issues raised by *pro forma* abstracts themselves.

The conduct of the study

As I have pointed out, this is only a preliminary and personal study. It lacks rigour and thoroughness in several respects. It needs to be followed up with a full-scale study, in which the classification of research is conducted by a panel of abstract-writers who have been properly instructed and trained, and have been given preliminary tests to ensure consistent results. The contrasting study of other fields of engineering needs to be widened, and the selection of fields made in a more objective manner. This preliminary study, while undoubtedly weakened by the use of informal methods, has perhaps indicated that there is a significant difference between HCI and other engineering disciplines, worth investigating further.

On the uses of *pro formas*

I have found *pro formas* useful in helping me discriminate more clearly between the different categories of research product. More than that, I have found that transforming the contents of papers into abstracts has involved me in reading and, I believe, understanding the contents much more thoroughly than I had before. I have learned through this process that the essential contribution of papers to system design is often far from obvious, at least from the abstracts provided by authors. This has led me to wonder whether it might be beneficial if authors were to include a *pro forma* abstract encapsulating the results of their research.

	EM	*ES*	*ET*	*RS*	*XH*	*Other*	*Tot.*	%*E..*
CHI '89	3	6	13	12	18	2	54	40.7
CHI '90	5	3	4	11	21	3	47	25.5
CHI '91	3	4	7	15	20	3	52	26.9
CHI '92	9	4	8	17	28	1	67	31.3
InterCHI '93	4	3	10	15	26	4	62	27.4
Totals	24	20	42	70	113	13	282	30.4
percent	8.5	7.1	14.9	24.8	40.1	4.6	100.0	

Table 3. Final results of analysis: distribution of HCI research papers among five categories, including Radical Solutions (*RS*) and Experience/Heuristics (*XH*). The "%*E..*" column shows the total percentage falling in the *EM, ES* and *ET* categories.

An advantage of writing *pro forma* abstracts describing one's own research is that it makes clear what the research does cover and what it does not. If no motivating analysis or confirmatory test has been conducted, these portions of the abstract are notably absent. Again, if authors were to attempt to cast their work into *pro formas* they might notice these omissions before it was too late.

The *pro formas* discussed here are intended to describe research after its completion, but they can easily be adapted to describe research projects before they begin or while they are in progress. They become useful for proposing research or reporting on progress, e.g.:

> A radical solution to the problem of interacting with the contents of paper documents will be developed, based on the use of an over-the-desk video camera and image recognition. In comparison with manual solutions or with solutions based on scanning, OCR and display, the proposed solution should offer much improved speeds of task performance, which will be demonstrated in preliminary tests. It can be expected to exhibit a number of side-effects including sensitivity to poor lighting and to paper movement, which will be investigated (based on [9]).

Again, an advantage of using *pro formas* to document proposed or ongoing research lies in their help in identifying essential components and activities, such as the essential comparison of proposed radical solutions with normal solutions.

Why invest in radical solutions?

A number of plausible reasons can be found for HCI researchers' apparent interest in radical solutions. The pace of technological change may very well encourage radical innovation by making existing solutions obsolete and new ones possible. However, this would not seem sufficient reason to persuade so many researchers to give up their traditionally analytical role to become inventors.

A possible explanation lies in the economic structure of engineering research, and the researcher's preference for the cheapest route to the biggest return [10]. When the enhancement of models is hard but the design and testing of solutions is easy, researchers are likely to gravitate the latter if they have the choice. HCI creates an extreme bias in favour of the designer, since the improvement of analytical models of human behaviour is so very difficult in comparison with the construction of prototypes for humans to try out.

But why engage in radical design rather than enhance existing solutions? Does HCI research perhaps offer a degree of freedom that other research engineers are denied? It is significant that many of the solutions proposed by HCI researchers are relatively unconstrained by the environment in which they are intended to be used; at least the constraints are much less than in Electronics, say, where there are numerous design standards to be complied with. There is an opportunity for the designer to give free rein to his or her creative abilities. Radical solutions are very likely to emerge if this is creativity is encouraged; which leads me to the second question.

Does industry need radical-solutions research?

We might expect to identify an industrial need for the products of any engineering research discipline. In the case of HCI, given its emphasis on design, we might expect to discover this need in the way design is done. The earlier account of the design process identified a number of ways in which the products of research can assist the designer; we would expect these to provide ways for radical solutions to make contributions.

If we look for opportunities for radical solutions to assist design, however, we see that they are inherently handicapped by virtue of their radical nature. For example, radical solutions cannot easily contribute to step 2 (choosing a solution) because of difficulty in comparing with other solutions, nor to step 3 (working out details) because their details tend to be ill-defined, nor to step 5 (estimating compliance with specifications) because they do not usually permit performance predictions. Far from assisting normal engineering design, then, radical solutions might be expected to hinder it, creating more work for the designer and increasing uncertainty and risk. We might expect to see evidence of design failure as a result of introducing radical solutions, and indeed we do see occasional reports of this kind, such as the investigation of the London Ambulance Service system failure and the report on Therac-25, both of which suggest that a more radical user interface design was used than might have been appropriate [7,8].

The point I wish to make is not, however, that radical solutions are harmful as research products. They are useful, indeed essential, in enabling system designers to break free from existing solutions that are reaching their limits of enhancement. The point is that the enhancement of solutions is vital to progress and is not being published. Radical innovation in fact breathes a great deal of vitality into HCI research and thus generates benefits of a special kind. It needs to be motivated by real evidence that existing solutions are reaching their limits. It needs to be complemented by a flow of enhanced solutions that designers can study, learn from and possibly exploit in their own work. This flow of enhanced designs is being stifled; we need to understand why, and do something to correct it.

ACKNOWLEDGEMENTS

I wish to thank my colleagures at Rank Xerox Research Centre and Xerox PARC for their helpful comments on this work. I am grateful in particular to Mik Lamming for his suggestion to use *pro formas* to describe proposed research. I am also most grateful to Peter Polson for his comments, and for his suggestions on experimental methods which I have summarized above. I wish to thank the anonymous reviewers, one of whom drew my attention to important prior work on research classification.

REFERENCES

1. Borning, A., and Travers, M. Two Approaches to Casual Interaction over Computer and Video Networks. *Proceedings of CHI '91 Human Factors in Computing Systems* (April 28-May 2, New Orleans, LA) ACM/SIGCHI, N.Y., 1991, pp. 13-19.
2. Card, S.K., Moran, T.P., and Newell, A. *The Psychology of Human Computer Interaction*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1983.
3. Constant, E.W., II. *The Origins of the Turbojet Revolution*. Johns Hopkins Univ. Press, Baltimore, MD, 1980.
4. Jarvinen, P. On Approaches in Information Systems Research. *Proceedings of 14th IRIS Conference*, Univ. of Umea, Sweden, 1992, pp. 99-112.
5. Johnson, W., Jellinek, H., Klotz, L., Rao, R., and Card, S., Bridging the Paper and Electronic Worlds: the Paper User Interface. *Proceedings of InterCHI '93 Human Factors in Computing Systems* (April 24-29, Amsterdam, Netherlands) ACM/SIGCHI, N.Y., pp. 507-512.
6. Kuhn, T.S. *The Structure of Scientific Revolutions*. Univ. of Chicago Press, 1962.
7. *Report on the Inquiry into the London Ambulance Service*. S. W. Thames Regional Health Authority, 40 Eastbourne Terrace, London W2 3QR, 1993.
8. Leveson, N.G., and Turner, C.S. An Investigation of the Therac-25 Accidents. *IEEE Computer 26*, 7 (July 1993), 18-41.
9. Probert, S.D., Marsden, J.P., and Holmes, T.W. *Experimentation for Students of Engineering, Vol. 1.*, Heinemann, 1971.
10. Newman, W.M., and Wellner, P. A Desk Supporting Computer-based Interaction with Paper Documents. *Proceedings of CHI '92 Human Factors in Computing Systems* (May 3-7, Monterey, CA) ACM/SIGCHI, N.Y., 1992, pp. 587-592.
11. G. F. C. Rogers (1983). *The Nature of Engineering.* Macmillan.
12. Rouse, W.B. *Design for Success*. Wiley, 1991.
13. Starker, B., and Bolt, R. A Gaze Responsive Self-disclosing Display. *Proceedings of CHI '90 Human Factors in Computing Systems* (April 1-5, Seattle, WA) ACM/SIGCHI, N.Y., 1990, pp. 3-9.
14. Vincenti, W.G. *What Engineers Know and How They Know It*. Johns Hopkins Univ. Press, Baltimore, MD, 1990.