The diversity and ethics of HCI

Paul Cairns & Harold Thimbleby
UCLIC, University College London Interaction Centre

This paper argues that diversity in HCI is legitimate, and indeed a source of richness. The diversity is directly analogous to the diversity in ethics, a field that has had millenia to develop and mature. This paper then addresses not HCI in society but, more specifically, the functioning and purpose of the HCI community as an ethically cohesive society. The paper exhibits some current problems, and provides some specific suggestions for progress.

Categories and Subject Descriptors: H.1.2 [**Models and principles**]: User/Machine Systems; K.4.0 [**Computers and Society**]: General; K.7.4 [**The Computing Profession**]: Professional Ethics

General Terms: Human Factors

Additional Key Words and Phrases: Ethics, Paradigm, Human Computer Interaction

> "Science will only maintain a 'human face' if there is a metaphysical system in place robust enough to provide a matrix within which intrinsic human value can be sustained." *Alexander* [2]

1. INTRODUCTION

Human computer interaction is a diverse field: how can computers, in all their forms, help computer users have a better experience? The effectiveness of users depends on many human factors, from how they learn to how they interact in social groups; and the effectiveness of computers depends on many factors, from how well their software is engineered to aspects of their physical ergonomics. HCI as a research field involves artists, mathematicians, and psychological experimentalists, to name a few. As a practical field, HCI involves industrialists and politicians. Moreover, even when consensus can be achieved across such diverse perspectives, the discipline is over-taken by technological developments, as well as market and consumer led pressure to deliver novelty.

In an established discipline, such as physics, diversity is at once both accepted and compartmentalised. There is theoretical physics through to engineering. Spe-

Address: UCLIC, University College London, 26 Bedford Way, LONDON, WC1H 0AP, UK.
URL: `http://www.uclic.ucl.ac.uk`

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works, requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept, ACM Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or `permissions@acm.org`.

cialist subareas, such as electronics, are further divided into power, analog, digital, microwave and so on. Workers in an area of physics draw upon the knowledge and practice of that particular area. In HCI, there is no such privilege — one aspect of usability cannot be neglected at the expense of another even, as we shall discuss, when conducting research. The resulting tension makes it difficult to progress the discipline in a scientific manner and, indeed, some HCI practioners are deliberately moving away from science and re-branding the subject as interaction design [32].

This paper discusses the source of tension and even confusion over what HCI is. We attribute the tension to a lack of working paradigm, to use Kuhn's term [17]. Furthermore, we contend that not only does HCI not currently have an appropriate paradigm, it will not have any paradigm in the foreseeable future. In Kuhn's view, this suggests that HCI cannot function as a science. We are not however claiming that HCI could or should function as a science like physics. Instead, in our discussion, paradigms are to be understood in the broader sense of the manner in which subjects organise themselves and function. Even in this broad sense, a useful HCI paradigm is elusive because of the lack of stable phenomena of study.

However, HCI does have a more established epistemolgical heritage, namely in the field of ethics. This provides the intellectual rigour and tradition to replace the narrow conception of needing 'a' paradigm. The long-accepted diversity in ethics corresponds, as indeed it must, to diversity in the field of HCI.

2. ETHICS

Ethics is often taken as a narrow professional ethics: in HCI, this would mean, for example, not exploiting users or taking due care when working with children. Most professional societies have codes of ethics; the ACM has one [1], and HCI practitioners who are members of the ACM special interest group in HCI (ACM SIGCHI) effectively submit to this code. The simple motivation for professional ethics is that it provides a generic framework to define 'best practice' that limits legal liability.

Burmeister [5] interprets the Australian Computer Society's code of ethics for usability engineers. His examples raise questions about how to respond to problems; for example, in a usability study of a web browser, a subject accidentally finds an objectionable web site — what do you do within the code of ethics? In contrast, in this paper we want to take ethics as a guide for positive action, not as a guide just for managing exceptions. As we shall see, this perspective itself mirrors one of our arguments about usability: is usability about removing problems, or is it about pursuing good?

HCI can be involved in the design of interactive systems that unequivocally raise ethical issues. For example, how should the user interface of a euthanasia assistant be designed? Should it have undo? Should it require more than one user? Somehow the rules written into the program that implements the assistant become moral rules imposed on (or unavoidable by) the users of the device; or, conversely, the moral dilemmas raised by the use of the device are either anticipated or ignored in the coding of the software.

Less obvious are the ethics of devices like mobile phones. The code inside a mobile phone imposes a way of use onto users, and for example influences how individuals can communicate using the short message service. Can errors be undone? The

Ethics	HCI
Deontology	Standards (e.g., ISO, NASA)
Situation ethics (ethical particularlism)	Depends on user's task, motivation, etc
Ethical monism	Usability
Utilitarianism	Cost-effectiveness of design
Consequentialism	Usability metrics
Virtue ethics	The designer is right
Hedonism	Promote user enjoyment

Fig. 1. Some example ethics/HCI correspondences.

tediousness of writing clearly, which is a direct consequence of the user interface design, has led to the widespread adoption of shorthand languages. Examples like "`CU L8R`" are at once a sign of user creativity, a symbol of being in a trendy culture, and a consequence of inefficient design — or perhaps a successful outcome of shoehorning a bigger application into a physical device that might otherwise have been thought too small to support it.

Being able to interpret HCI ethically is not so surprising. HCI is a normative science that aims to improve usability. The three conventional normative sciences are æsthetics, which deals with things that are admirable, ethics, which deals with things that are good, and logic which deals with things that are true [30]. Broadly, HCI's approaches can be separated into these categories: logic corresponds to formal methods in HCI and computer science issues; modern approaches, such as persuasive interfaces and emotional impact, are æsthetics; and the core body of HCI corresponds with ethics. Some might add to this list rhetoric (using communication effectively), but, whatever terms we use, HCI is about making the user experience *good*. When considered further, the parallels between ethics and HCI are substantial (see Figure 1). This, then, is no accident. It would be informative to explore the relationships deeper but this paper is not the place to review the vast subject of ethics *per se*! Instead we exemplify the value of ethical parallels in HCI by taking just three particular ethical stances and interpreting (some) HCI work from these positions: Christian ethics, medical ethics and Rawlsian justice.

Regardless of one's commitment, indifference or opposition to these particular stances we have chosen to discuss, they provide coherent ethics. Of course there is yet more to the ethical implications that we do not have space to debate here, just as there is yet more to HCI itself: HCI for game playing, learning environments, persuasive applications ('captology' [12]) or secure systems (e.g., voting) highlight some further complementary ethical issues that are not explored here. The point is, for the purposes of this paper, that drawing out specific ethical relations with HCI is both challenging and rewarding; using ethics *constructively* raises relevant and perceptive issues in HCI — and ones that practitioners may make either implicit or explicit commitments to. Further, if practitioners make ethical commitments to an established framework, debate about that choice becomes routine.

2.1 Christian ethics

HCI has a set of basic assumptions. We should empower users. But why? Why not exploit users and maximise corporate profit? We should improve usability. But

why? Why not encourage users to buy upgrades, and encourage market churn? Users make mistakes, so user interfaces should have undo. But why not make users live with their mistakes?

Clearly, HCI — or a different HCI — could work from a different set of assumptions.

Christianity takes a set of assumptions from the Bible. Humans are diverse and made in the image of God; because people are valuable, it is worth striving to improve their circumstances; but creation is fallen, through the agency of sin, and attempts to improve situations are likely to fail unpredictably. Christians deal with this by promoting grace and forgiveness. And so on ... (this paper is not the place to pursue any particular framework in depth; the technical Christian terms can be followed up in [16]).

Compare these ethical positions with those of HCI. The basic assumptions of HCI are that users are diverse and valuable. It is worth developing new systems that improve users' circumstances. However, the world is complex and designers build systems for them that then fail in unexpected ways. These shortcomings can be ameliorated by things like better user training, or better error recovery — but these may also fail in interesting ways. Iterative design is recognition of the fallen nature of systems. That is, the current system is not as it should be: as good HCI practitioners, we should continually strive to improve it.

Christianity and science do not currently sit comfortably together. Historically, however, the Western development of science (which contrasts with Chinese and Arabic developments, and indeed with developments under any other worldview) was motivated by the monotheistic imperative that nature has a structure and purpose. Even Einstein, who did not believe in a personal God, held that "God does not play dice" [11]. Many of today's ethical dilemmas (e.g., euthanasia, abortion, cloning, genetic modification, industrial pollution) arise because of an often unspoken tension between 'raw' technical developments and implicit Christian ethical habits of thought. Being clear that the HCI ethical orientation emerges from the Western/Christian matrix highlights responsibilities — a Christian might put it thus: good HCI is not optional because users are sacred. The answer to the lawyer's famous question to Jesus [22] is, "the user is your neighbour."

2.2 Medical ethics

Medicine shares a lot with HCI [40]. People have medical problems and doctors want to fix those problems. There is a whole range of medical approaches. Physiotherapy trains people to cope better — like HCI's emphasis on training. There is corrective surgery that like HCI might do "hip replacements" to change one system for a better one. There is preventive medicine, health, fitness, forensic, pharmaceutics and other branches. Both medicine and HCI are trying to improve people's lives and both have the knowledge and power to do so. If you invent a drug to combat a disease you transform the world; if you develop a user interface (particularly a web-delivered one) you transform the world.

Medicine has been around a lot longer than HCI, yet it has survived in a recognisable form since the Hippocratic Oath was created. Indeed, the Oath constitutes an ethical stance on what medicine is and how it should be practiced. The Oath, although now obsolete, brought coherence to the discipline in terms of what is and

is not acceptable practice. Though medicine has changed radically over the millenia and new technology (the same technology driving HCI) is changing it at an ever-increasing rate, there are clear and exacting procedures that must be followed if you are claiming to do medical research. New drugs are only introduced after a long, staged series of trials; trials must be of a certain size and controlled; reports of results must allow other clinicians to repeat the work. Different countries have different rules. And so on ...

With Christian ethics we illustrated an implicit connection that does not not necessarily determine detailed procedures (e.g., controls or randomised trials); but with medical ethics we illustrated a more prescriptive ethical framework, which if taken seriously would challenge HCI's *de facto* ethical standards. HCI would do well to learn from the processes that medical ethics requires. When we develop new web technologies, affecting millions of people, are the experiments done carefully, or with half a dozen students? Do we only do experiments on our own familiar systems? Are we satisfied announcing results with only three out of five patients cured? Do we use placebos, double blinds, controls? Did our published output provide enough details or enough of a working system for other people to check the results?

At the very least, HCI researchers should work to basic clinical standards when we do get to help design medical systems. If further we are to improve HCI more widely, one of the places to start is in our procedures for publishing and disseminating best practice. In medicine, any reliable knowledge is worth building on; in HCI practice we are more often concerned with the excitement and business side of things than the quality of the science and the ethics of doing it.

Thus, the medical ethical stance has a lot of mileage when adopted in HCI. After all, both fields are improving people's lives, but medicine has been doing it longer, so we surely have much to learn from it.

2.3 Rawlsian justice

A standard problem in ethics is to operationalise it; how and under what conditions do you *do* good? Rawls provides a thought experiment in the specific area of justice [34]. Rawls imagined a political world with rules. In what sense can this world be just, and if so, how can the rules be defined?

First, note the analogy with HCI. In HCI we are concerned with creating 'worlds' for users that are just. A device like a mobile phone, say, imposes a set of rules of interaction on the people who use it. In some sense the HCI designer wishes to bring about a world in which justice is increased. Rawls, of course, considered this problem in the political domain, as in how would a society be brought about where there is an increase in justice.

A key problem is that politicians are selfish, and unlikely to define just laws. Rawls's solution was to envisage a *veil of ignorance*. Politicians would design the society in ignorance of what future role they would play in that society. Only if a politician did not know whether they would be rich or poor, single or married, disabled, young, isolated ... would they devise just laws.

In HCI, we do not know what sort of users our systems will have. One of the standard slogans in HCI is that, unless designers consciously do otherwise, systems are accidentally built for their designers, not for their users. In a sense, unless

designers consciously work under a Rawlsian veil of ignorance, they are susceptible to building systems they like, rather than systems that benefit their users.

If a designer considered that, in the world where their system is imposed, they might be working on a user support line, or they might be writing the system's user manuals, or they might be actual users — then they would design with more consideration. In fact, for any successful system, designers are more likely to be ordinary users than any other category of user. Steve Job's famous exhortation to reduce the boot time of early Macintosh computers is a case of this insight : every second a Mac takes to boot, there are so many users waiting, that cummulatively whole lives are being wasted. On Rawls' account, developing user interfaces to make weapons of mass destruction easier to use would be unjust, since no designer would wish live in the target population.

The two contrasting examples, Mac user interfaces and weapons interfaces, shows how a Rawlsian view of justice leads to HCI professionals probing and then making ethical decisions that influence their activities. Further discussion and a worked example in HCI can be found in [10].

2.4 Summary

To summarise, the three ethical stances addressed here give concrete and useful insights into HCI practice. Christian ethics already has a strong influence on current HCI values; medical ethics has similarities but more exacting standards; and Rawlsian justice is already partially, albeit implicitly, employed as a design maxim. Even if HCI did not require an ethical meta-perspective, explicitly describing and developing HCI via these stances would bring new and *good* practices to HCI. However, as the next section discusses, HCI requires a meta-framework where values can be debated because it is incapable of having the inherent structure of a normal science.

3. HCI EXPANSION

Computers are now pervasive, and new and innovative uses are being found daily. HCI has changed and adapted from its early beginnings in microcomputers and the organisational issues of the 1960s timeshared computer systems to the personal and international internet needs of the 2000s: a vast change, both smaller and bigger. HCI is now one of the most rapidly growing areas in computer science [35]. In such a flourishing environment, it is surprising that HCI is not gaining strength and momentum. Instead, it seems fraught with re-branding as interaction design [32], new usability [44] or funology [23]. These changes of nomenclature reflect deeper disagreements as to what HCI is, disagreements reflected in the apparent absurdity of reviews of papers and research proposals. The authors of this paper have had different referees say *of the same work* that it is internationally outstanding research, and not worth an undergraduate diploma.[1] In essence though, it seems that these features in HCI are symptomatic of a single cause: HCI does not have a paradigm

[1]Note: we are not just authors; we are also referees ourselves, and we admit "from the other side of the fence" that it is extremely difficult to referee HCI (and HCI-related topics) fairly. That is, we are not complaining about referees, as if authors are right and referees wrong; we are referees too. Both authors and referees are part of the same community.

[17]. Some evidence for this is laid out below. Straightforward analysis shows that the reason for the lack of paradigm can be simply attributed to consequences of Moore's Law [24]. This has deeper implications: not only does HCI not currently have a paradigm, but that whilst Moore's Law holds it will be impossible for HCI to have a paradigm.

3.1 Paradigm Lost

Paradigm is a now well-established term in the philosophy of science as a result of the work of Kuhn [17] and is broadly adopted as a tool for investigating social phemonena, too. (For example, the impact on business of 'disruptive technologies' [9] is analogous to Kuhn's crises.) According to Kuhn, when an area of science has a paradigm, there is some fundamental set of assumptions agreed upon by that particular scientific community. The full extent of these assumptions are rarely made explicit but allow the science to proceed "normally." That is, the paradigm has some descriptive and explanatory power, and normal scientific research works to extend the paradigm to cover previously unexplained phenomena. Furthermore, the questions asked whilst doing normal science are usually obviously answerable, at least in principle, in terms of the paradigm — Kuhn saw normal science as routine puzzle solving. The novelty of the research exists in working out details. As Kuhn describes it, there is usually a particular paradigm embodied in a theory used to conduct research in the area. Kuhn's historical perspective of how science progresses is not universally accepted (e.g., it contrasts with a Popperian account [13]), but paradigms remain a central and perceptive concept for describing and distinguishing normal, incremental and coherent scientific research by a community of people.

It is easy to point to theories in HCI that have enjoyed success: GOMS [6]; affordances [25]; situated actions [37], say. Like paradigmatic theories, to a degree, they underlie a lot of current practice in HCI. However, a scan of important HCI journals shows that there seems to be little normal science being done to extend or refine these ideas to new phenomena. Indeed, Norman is even rebuking HCI for the misuse of the term 'affordance' to describe certain interactions [26]. So, though these are relevant and valid theories, they lack the power and focus of paradigms to aid a coherent research effort.

Although it helps to give specific examples, the flaws and strengths of any individual theory (such as those mentioned above) are not relevant here. Nonetheless, leading researchers in HCI declare themselves to be still looking for theories of HCI. Shneiderman has said that he is "... looking for a generative theory of HCI" [Personal communication], Peter Johnson is "trying to produce a theoretical framework for interaction. As are we all." [Personal communication] and even we want "a way to analyse [user interface] designs before they are built" [42].

The consequences of an absence of theory are also evident. Kuhn describes some features of a 'pre-paradigm' or 'between-paradigm' science:

—Squabbles over what are legitimate tools for research
—Squabbles over what are legitimate phenomena to study
—Inability to scope the domain of study

Evidence for some of these features can be found in Höök's discussion list in CHI

Place [15], the purpose of which was simply as a forum for discussing the referees' comments on recent CHI 2003 submissions. In particular the thread initiator asked for referees' comments that seemed particularly irrelevant or ludicrous. Assuming that HCI is not filled with mischievous or incompetent referees, ludicrous or irrelevant reviews could only come from the inability to communicate. That failure must also be in both directions unless HCI has an unusually large number of mischievous or incompetent authors failing to understand reviewers. A plausible argument for this failure to communicate would be an absence of shared knowledge and framework in which to argue. Or to put it another way, a lack of common paradigm.

A more surprising contribution to Höök's list was from Jonathan Grudin. As an editor of this journal, as well as a researcher of some experience and standing, he surely knows what constitutes a good HCI article. However, he too bemoaned that he had recently had rejected his "best CHI submission in a decade." This may be a common occurrence in HCI but in mathematics, say, it would be unthinkable that a significant researcher (not yet in their dotage) would be rejected at a conference. This is not because mathematicians always produce publishable research but rather because any researcher of sufficient experience and maturity would only offer good things for publication. And if they were mistaken and their research was flawed they certainly would not get in a position of lamenting unfair treatment. Grudin's case, then, seems again to be an instance of referees and authors failing to engage on a common theoretical basis even when the author has good authority for the validity of his research.

As Kuhn points out though, the absence of paradigm need not be a cause for concern. It is a legitimate, recognisable phase in the growth or emergence of any science. Through argument over legitimate objects and tools of study, new theories emerge with greater explicative and predictive power. Comments from a number of contributors to Chalmers's list [8] indicate that they are expecting a theory or theories to emerge. We argue that this is unlikely to happen in the near future.

3.2 Phenomenology in HCI

In the natural sciences, any theory is aiming to describe a set of phenomena. Quite which phenomena are targeted by any particular theory tends to be built into a paradigm rather than chosen on a more logical basis. Gravity, the mutual attraction of all massive bodies, was not studied as a phenomenon until Newton described it. Once a particular phenomenon has been observed and recognised as important, successive theories must either explain it or incorporate it into the theory. For example, when faced with the problem of the constancy of the speed of light in Newtonian mechanics, any new theory would have to predict the constancy or build in the constancy of the speed of light, as did special relativity [11]. Once there is a theory, new ideas are highlighted that nobody was aware of prior to the formulation of the theory; indeed, coincidences or anomalies that may not previously have been noticed — or, indeed, that people refused to notice because there was 'no' compelling explanation — may now be explained in the terms of the theory [20]. In HCI we do not have the conceptual structures provided by compelling explanations; there are few if any anomalies begging for explanation. Often problems are trivialised because the theory is deemed inapplicable in the troublesome area — none of the

theories available to HCI claim to have the generality and precision of, say, gravity that affects *all* bodies in the *same* way.

Furthermore in HCI, the phenomena of study are much less enduring, simply because they arise as a result of interaction with an artefact. Artefacts are not externally given but can be modified by their makers. Thus the phenomena of HCI are not givens. For instance, a theory may explain the reasons why paragraph formatting is difficult in a particular word processing package, only to find that the next version does not use the same method and the whole problem has gone away. And even if this did not happen, another argument might be to say, "Well, if that package is problematic, use this one instead."

Of course, if the artefacts of study of HCI were essentially static, such phenomena may vary between packages but a particular theory of say, how people interact with word processors, could eventually be possible and proscribe and prescribe particular interactions (or whatever). In reality, of course, HCI's artefacts are far from static.

It is worth pressing this comparision further. In biology, we recognised a huge diversity of life. Its classification is a research field in its own right (e.g., phylogenetic systematics), and is one paradigmatically recognised as worthwhile. It is fair to say that the diversity of HCI artefacts is much less than the diversity of life, yet a cladistics of HCI would be pointless: there is no sense in principle that it could be completed, and none that a completed classification, even if it were possible, would be of value the following year.

Although first put forward in 1965 and anticipating phenomenal growth in microelectronics, Moore's Law is now taken to describe and predict a doubling of computing power (e.g., calculations per second) every eighteen months. Thus, at a purely algorithmic level [14], computers now can do about 100 times as many calculations per second as computers ten years ago. Moreover, that statement will still be true in ten years time. Many people, however, predict that soon Moore's law will hit the physical limits of silicon and so will no longer hold. Brown and Duguid argue that Moore's Law is not a physical law but an economic one [4] — more akin to a measure of competition and success in chip manufacture than ability to manipulate materials. Thus, Moore's Law is likely to hold for the foreseeable future.

In terms of HCI, Moore's Law may seem to have little bearing. After all, humans are not changing at an exponentially increasing rate and Moore's Law only relates to raw computing power. Consider then Moore's Law as interpreted in terms of screen images. It means that for the same size visual display, a computer can do 100 times more work to generate a single image than ten years ago. This makes possible virtual environments, gaming styles and visualisations that were not conceivable without that scaling up of power. Accordingly, the richness of potential human computer interaction is also scaled up.

Visual display is only one area of interaction that could be potentially improved by Moore's Law, but it demonstrates its implications. In parallel with Moore's law, there have been astonishing reductions in the cost of memory chips, huge increases in the capactities of hard disks, massive reduction in power usage and screen displays, enormous increases in network connectivity. All these things have acted to move computers away from a traditional Turing machine model of computation to more dynamic, interactive computation [45] where computers and humans all act

as a collaborating network of agents. In practical terms, legitimate interest of HCI in areas of ubiquitous computing, pervasive computing, implants and so on is now feasible — and challenging. Because of the rate of innovation, users have become informal HCI experts — exercising significant HCI preferences in their consumer choices, in turn driving new development; conversely, HCI professionals have become users. (Contrast this accepted, almost unavoidable position, in HCI with the demands of medical research ethics.)

In trying to produce a paradigm or core theory of HCI, researchers need to define and scope HCI phenomenologically. Or rather, producing a paradigm will entail an appropriate set of phenomena for HCI to study. But whilst Moore's Law and its counterparts in networks, storage media and miniaturisation continue to hold, any circumscribed set of reified phenomena will be rapidly swamped in the exponentially expanding potential of computing power. Even if a phenomenology for interaction once existed, it certainly does not now. The implication for a paradigm or theory is that not only is there currently no paradigm but, in the persistent increase in available computational power, there can be no adequate paradigm.

4. EXPELLED FROM PARADIGMS

Without some meta-framework, such as that which a paradigm might have provided, we claim the expansion of HCI phenomena will always be too much for any individual to grasp. The result will be that HCI must become a loosely held community of "islands" where each island has its own acceptable basis for research but which has no easy communication with other islands.

Without a paradigm, HCI cannot function as a normal science, in Kuhn's sense. This is not alarming — there are many disciplines that make no claim to being natural sciences and yet which function perfectly well. HCI could perhaps borrow a framework from some of these other disciplines.

The simplest of these would be to fall back and say that HCI is entirely a craft discipline. This has been explored extensively elsewhere [21] and it is not an appealing position, for then where is the science of any sort; where are the aspirations to truth? How can research be anything other than *ad hoc*?

Design subjects, such as architecture or graphic design, have an obvious relevance to HCI. Yet, there is no more a theory of architecture than there is a theory of HCI. This is not to say that design subjects are without a paradigmatic foundation. At the very least, in all but electronic media, there are concrete phenomenological constraints. Buildings are made from physical materials. Graphics are limited by the manufacture and reproduction quality.

In addition, there is a persistence to non-virtual designs absent in computer-based designs. Architects in the sixties raved about the innovation in high-rise inner city developments and the opportunity for new types of communities. They did not then immediately pull them down to build Version 2 of the same buildings. Instead, the buildings persisted, the communities grew into them and the grave flaws with the designs became apparent. And architects learned from them [19]. This sort of learning from past designs is an essential part of any education in a design subject.

Not only does HCI not have persistent phenomena, but analysis of existing designs, such as Paint Shop Pro or MS Word, are regarded as merely product reviews.

This point of view is understandable because, after all, no-one is going to have to worry about these designs in couple of years time. As a more substantial example, the Canon Cat is now regarded as an excellent piece of interaction design [33]. Many older people hold it as the epitome of good design — and that computers have taken a lamentable down turn since then. But only older people. Young students doing HCI degrees now have never even seen a Cat (some weren't even born when it was being produced!) and are never likely to do so either. They cannot re-perceive the Cat for themselves and recontextualise it to novel designs. They are reliant on (for instance) Raskin's description of it and he can only tell them his own perceptions. The Cat can no longer be demonstrated widely (if in fact any working models of it still exist); and exhibitions or pictures of it do not capture or illustrate any HCI issues beyond its unexceptional physical form.

Thus, to call HCI a design subject is attractive but misleading. HCI does not share a basic phenomonelogical reality with other design subjects. The community of HCI professionals cannot build on each others' work and a long-standing shared heritage of exemplars, about which 'everybody' has paradigmatic opinions.

HCI could instead be regarded as a social science. Muller specifically makes this link in criticising the need for a theory by refering to the absence of theory in Woman's Studies or African Studies [8]. But once again, there are concrete phenomena underlying social sciences — people. Admittedly people are immensely complex in their social interactions but: they are not evolving (not, at least, at a rate significant with respect to the advacement of knowledge); individuals do not double their performance every eighteen months; societies can only really transform entirely on the scale of individual longevity.

This may seem like an opening for HCI — it is really a social science about people and how they use computers. Social interaction is certainly a fascinating aspect of HCI but if it were only that then we would be abdicating from trying to influence the computer side of the discipline. In effect, to view HCI as only social science is to fail the millions of users who suffer from poor interaction design everyday.

4.1 Summary

What this tells us is that HCI, like all disciplines, must carve its own framework in which to conduct research, develop skills and accumulate knowledge. The allure of a natural science or social science or design subject framework is sadly superficial. However, by thinking about HCI as a normative and therefore as an ethical science, there is a possibility to adopt from ethics the appropriate framework for the study of good user interfaces.

5. PRACTICAL SOLUTIONS

We have already illustrated how particular ethical stances can help interpret and stimulate work in HCI. Clearly in this paper, to explore all areas of HCI or ethics and map out a full correspondence is surely to fall short of what is necessary. Instead, we draw out some practical solutions to problems in HCI as a discipline. These are based on how ethics, as a discipline, is able to solve its own problems.

5.1 Privatives

For any ethical position to influence, rather the merely comment on, behaviour, it needs to be operationalised like Rawlsian justice discussed earlier. If HCI is taken to be about making good user interfaces, what constitutes a good interface becomes very difficult to define. When "good" is not operationalised, it becomes very difficult to assess the value of work.

As one concrete example (though there are potentially hundreds), in a classic HCI paper, Young [47] describes how surrogates and task-action mappings make user interfaces more usable. In particular, Young provides a detailed critique of hand-held calculators, and the user mapping their task into operations to command the calculator to achieve tasks. This analysis exposes some interesting design issues for calculators and similar styles of user interface, and contributes to the understanding of usability in such domains. One might think that the analysis therefore improves usability. If so it should be feasible to take an alternative domain, such as virtual reality or even a radically different calculator [39], and be able to apply the method. This turns out to be very problematic. In which case, the method competes with other conceptions of usability. For example, it is clear that task-action mappings ignore motivation, social context, error recovery and numerous other perspectives on usability. One might reasonably criticise task-action mapping as inadequate to achieve usability.

As a move towards a form of usability that can be worked with (rather than one that encourages picking problems), we propose that usability is best thought of as a *privative*.

Since Aristotle, a privative concept is a negative category of thought. In everyday English, a 'privation' is an absence of something wanted; the word privacy, for example, implies absence of public benefit. In many areas, the privative and positive are equivalent for practical reasoning. Another example: cold is the privative of the positive hot. We might draw curtains to keep out the cold, and this — as its very common usage suggests — makes sense. We put things in a fridge to cool them down. We put on jumpers so we don't get cold.

Is cold just the opposite of hot? Unfortunately, coldness is not quite the opposite of heat, for there are some cases where it does not work as naïve reasoning might lead one to think. For example, a heater with a parabolic reflector can heat an object at a distance; on the other hand, a 'cooler' with a parabolic reflector does not work. The reason is simple: heat is radiated *from* a hot object; even some heat is radiated *from* a cool object. Once there is a clear concept of energy (e.g., electromagnetic radiation) it is clear that heat does work, and the idea of cold 'doing' things is an informal, albeit convenient fiction.

There are many such historical examples of showing how science has progressed through privative concepts until appropriate theories ruled in favour of a positive version. Most famously, phlogiston was a very successful theory until precise measurement of mass and the discovery of oxygen suggested that phlogiston was better thought of as the private version of oxygen rather than treating it as the single negatively-massed element. Conversely, the discovery of electrons as charge carriers made no difference to Benjamin Franklin's attribution of current flowing from positive to negative because we had no prior view that electricity had to be positive

(as we do for mass).

More recent examples of the value of privatives for explanation (so-called *narrativium*) can be found in the alternative science of Discworld, the popular fantasy milieu created by Terry Pratchett [31]. Our analogy here is not purely made in jest! Pratchett creates an alternative world, much like HCI creates new worlds for users — indeed, worlds that envisioned only a few years ago would have seemed to be science fiction too. Pratchett argues that worlds are held together by their narrativium, the narrative sense they make to readers — or in HCI, the sense they make to users.

Now to the point. Usability can be understood as a privative. Usability is the absence of usability problems. Problems are the positive concept, and problems are what designers and users have control of. To the extent that phlogiston, cold, and other privatives seem to make sense, there is some sense in 'making' systems usable. But this sense disappears when one tries to explain rigorously what in fact is happening. Narrowing usability to a less interactive more ergonomic sense, a system is usable if it is comfortable but in ergonomics, comfort has already been recognised as a privative concept [29].

With usability viewed as a privative, Young [*op cit*] identified one way to analyse a failing in usability, namely that of awkward task-action mappings, though expressed in positive terms. Thus, Young's method cannot be expected to improve usability universally — only to remove certain types of usability problems.

Crucially, seen as a privative, it is now legitimate to have alternative ways of critiquing usability. Motivation, for instance, which task-action mappings ignore, is no longer an *omission* of Young's position, but now is an alternative *complementary* framework to highlight other sorts of usability problem.

5.2 Complementarity

The rich diversity of HCI leads to conflict. What is the right way to do HCI? Our right way need not be the same as your right way — taken to its limit, this makes HCI all art and no science. If it is possible that HCI is big enough to accommodate more than one right way, then an exclusive view of what is right will exclude effective approaches. For example, a research paper on a new user interface submitted to a conference might be rejected because it did not do usability studies; or a usability study might be rejected because it was a laboratory study and not a contextual study; and a contextual study might be rejected because it studied a proprietary system. And so on ...

Niels Bohr introduced the idea of *complementarity* so that physics could satisfactorily handle the superficially contradictory ideas of, for instance, light as a wave versus light as a particle. Waves and particles are different levels of explanation, that are deeply equivalent, yet apparently incoherent: at our human-level of experience, nothing can be both particle and wave. Complementarity has also been used for the apparently different approaches of science and religion: both can talk about the same phenomena, such as rainbows, but they provide complementary descriptions, one in terms of structure and one in terms of purpose: how and why [2].

In HCI we have many complementary descriptions that superficially seem to contradict each other. How can we consider issues in HCI unless we consider user

motivation? How can we consider issues in HCI unless we consider the formal specification of the user interface? How can we consider issues in HCI unless we do sound empirical work? How can empirical work be sound if we do not have a formal definition of the system, whether the user or the computer? Given what we know of human diversity, how can it be sound without being an impractically large experiment? And so on ...

Whilst the privative sense of usability allows us to say *what* a piece of work is achieving, it does not necessarily motivate *why* it is worth achieving. However, interpreting HCI in an ethical stance and making that ethical stance explicit provides a 'why.' Thus work from a deontological stance should be seen to make a contribution to standards in HCI. A reader of such work may question the deontological stance but *within it* there are clear grounds (or there should be) for understanding the contribution of the work. Or again, a piece of work in formal descriptions of user interfaces is clearly in the logical stance and should be interpreted as such. Moreover, it is not worth stating at the same level of description that logical work makes no contribution to æsthetics or user motivation — these are not the concerns of logic.

5.3 Summary

Complementarity in HCI can be understood as the balance of work done in HCI from different ethical stances. Any stance defines and motivates what the work does and does not do, much as a paradigm does in a natural science. Now, with complementarity accepted, HCI is not committed to an untenable or fast-retreating paradigm, but is able to grow within different ethical stances or through comparison and differentiation of ethical stances. One clearly separates the two modes of work.

6. HCI IN THE WORLD

Perry's classic account of ethical development [27] shows how students (for instance, HCI students — us as we once were) generally start from an absolutist position: "There is one right way to do HCI." This initial position matures through uncertainty, relativism, and then through stages of personal ownership and reflection. At the highest levels, Perry argues a student makes a personal commitment to the particular ethical framework they have chosen to undertake their work. However, it is a dynamic activity to develop any such framework in the complex and conflicting real world.

Particularly with mobile devices and the world wide web, HCI research is impacting and changing humanity faster than many other fields. It is arguable that basic health and agriculture and infrastructure (e.g., potable water) are of greater direct benefit, but behind the success of any programme is literacy, and the usability of computers (e.g., through digital libraries [46]) are certainly a great part of our hope for humanity. In this perspective, HCI plays on the field of world ethics. (Indeed, cultural issues is rightly an important specialism of HCI.)

Küng has put it like this, "an answer in negative terms can hardly be enough if ethics is not to degenerate into a technique for repairing defects and weaknesses. So we must take the trouble to give a positive answer to the question of a world ethic." There are echoes of our comments on privatives, but we do not need to paraphrase Küng's programme here [18]; instead, we pick up one of his specific proposals.

In any situation it is difficult to weigh up benefits, and to balance individual, institutional and professional ethics. Küng paraphrases some rules from 'present day ethics' from his international and cross-cultural survey; whereas Küng was thinking of progress in general (e.g., gene manipulation, nuclear power) we add an HCI gloss to each one:

A rule for solving problems There must be no progress that brings greater problems. This is practically the wake up call for HCI: there must be no system implementation that brings about greater problems in its use.

A rule for burden of proof Anyone who presents a new system has to demonstrate that what is embarked on does not cause human or ecological damange. What grounds does the HCI research or results justifies the innovation?

A rule for common good Küng's example for this principle is to have preference for preventive medicine rather than remedial medicine. In HCI, the preference might be for easier to use systems than more intelligent help systems — which in HCI might merely increase the 'acceptable' levels of complexity the user has to cope with.

A rule of urgency The more urgent value (e.g., survival) has priority over higher values (e.g., privacy). Thus concern for a user's RSI (repetitive strain injury) or upper limb disorders could take priority over privacy considerations that arise in monitoring their work load.

An ecological rule This rule is a special case of the previous: survival of the ecosystem, which cannot be replaced, has priority over social systems. In HCI, we should consider the wastage that disposing of obsolete user interfaces causes [43]: we are usually so excited about the improvements in technology that we forget that to benefit from them we have to discard old systems.

A rule of reversibility Küng explicitly says in technical developments, reversible ones have priority over irreversible. In HCI this rule could apply at the micro-level (user interfaces should have undo), to larger scales, as in a system introduced should not make the users (or customers) irreversibly dependent on it.

7. WORKING IN THE HCI COMMUNITY

As HCI practitioners or researchers, we are embedded in the community and what HCI is or will become is determined by the consensus, and the consensus is balanced against personal inclinations and preoccupations.

The opinions of everyone in the HCI community matter. As an academic field, HCI advances when referees agree that ideas can be published in conferences or journals, or agree that ideas can be funded and studied. Ideas that are published or funded lead to views of success and also of competition. Naturally, not all ideas can be published and not all can be funded, but there is an important difference. Strict refereeing for publication increases quality by providing appropriate feedback to authors and ensuring a consensual standard of publication. On the other hand, whilst strict refereeing for funding should have the same effect, any negative criticism encourages funding bodies to fund other subjects. Funding bodies cannot tell the difference between bad work and good work that could be improved.

Unfortunately as funding gets harder, competition for the resources increases, and referees tend to become stricter in applying the standards of their own view of good

HCI, namely by using their ethical commitment. Clearly this is counter-productive to HCI as a whole, and in the long run also to the sub-discipline. Thomas Gresham originated the well-known Gresham's Law, "bad money drives out good money"; perhaps today he might have said, "bad HCI drives out good HCI." The root of the problem seems to be the expansion of what are legitimate phenomena for HCI to study.

The same happens in publishing HCI work, but less dramatically than funding HCI research. The resource is not cash, but attention. So much stuff is done in HCI, that referees filter out work, and what is left for publication is perceived as *the* good work. Worse one may then associate the dominant work that is funded or published as the only valid approach.

Whether an established academic gets funding or gets a paper published is in itself a minor concern compared to two other issues. Contract researchers may lose their jobs on the arbitrariness of referees. PhD students may waste three or more years using methods their examiners do not appreciate. In fact, PhD students are perhaps worst off, since the field is developing so rapidly that a sensible research decision made three years ago can be very hard to justify at the time of submission.

With such a rapidly expanding phenomenology there are many incompatible ways to evaluate ideas. What might seem like a very promising idea for a psychologist might be retrograde for a technologist (it might be based on now-obsolete standards, for instance). Indeed, because HCI is so big, whatever the proposer plans, there is always a much better way of doing part of the research that the referee can spot; there is always something the proposer has not read or does not know. Because HCI is also relevant in business, there are always cutting-edge products that in some senses do better than anything that can be done by individual researchers. And so on ...

Conversely, the referee may well not be an expert in the proposer's area but is still expected to make a sensible judgement. Here, the waters are muddied because HCI is closely tied to consumerism. We all buy and use interactive devices, and we are therefore all entangled with the subject itself. Our personal preference for, say, Macs over PCs inevitably influences our HCI knowledge and hence influences what we think is appropriate for HCI research. One person might think an HCI problem needs studying in a laboratory; another might think the solution is to be found in an available commercial product. We tend to want exciting product upgrades in our own lives, rather than deeper analysis of problems we think we understand; do we also want exciting research rather than deeper analysis? Unfortunately commercial products illustrate ideas without reflection or generality. The ideas are not carefully defined, and, because of that, they aren't scientific [41]. Indeed, standard and uncontentious scientific criteria (e.g., the medical ethics we discussed) also seem at odds with the criteria-in-use in HCI. Of course, there is a valid complementary theme in HCI of envisionment and fun that need not be supposed to be scientific ...

Many of our HCI standards are derived from current commercial products and demonstrations. One of the problems of HCI is that just reading newspapers or watching TV gives one the strong impression of being at the cutting edge of research. We are all armchair experts in HCI! Yet what is reported is not deep; it trains us to want new products rather than more understanding.

Ethical complementarity provides at least some framework for the basis of eval-

uating each others' work, be it for publication or funding. A referee may criticise work but if the criticism is about the ethical stance then that is distinct from criticism about the work within some ethical stance. For example, HCI developed in business may have a strong utilitarian aspect. A referee may criticise the work either for not being good enough in the world of utilitarian HCI or for being utilitarian. But these are distinct categories of criticism. A funding body may decide that utilitarian research is acceptable and so fund the work despite the fair criticism that it does not address (for instance) the needs of minority user groups.

Also, the ethical stance of both authors and referees allow them to communicate and understand each other. Hopefully, with an explicit set of frameworks in which to work, the days of seemingly "irrelevant or ludicrous" referees' comments will soon be numbered.

8. CONCLUSIONS

HCI is rooted in the convergence of computers and humans. It necessarily has aspects that are formal and mathematical, from the computer side, that underpin the engineering of systems; and it necessarily has aspects that are social and human, that underpin understanding how humans interact, with themselves, with computers, and through computers. Some parts of HCI are scientific, theory-based; some parts are craft-based; some parts are æsthetic. HCI has an economic impact, not just because better HCI makes people more productive, but because user interfaces are embedded in consumer products. Some products survive because of market forces rather than good usability. Some parts of HCI prosper because, frankly, we like to possess and play with certain sorts of systems.

Like 'usability as privative,' then, the claim is that HCI debates can be lifted from argument from different premisses to either meta-HCI (i.e., ethical) debate, or as appropriate debate from shared approaches. To express the same thought less emotively outside of HCI: a Marxist and an environmentalist will almost certainly disagree; they might do better to have a meta-ethical debate. Of course, debate itself is a process that can lead to good outcomes, and this is within the arena of discourse ethics. The point is that a colleague in HCI is not necessarily wrong when they argue about usability from a different point of view. The positive view of usability encourages a monist approach to the field and tends to encourage categorising alternatives as ineffective. Viewing usability as a privative is a simple step towards clearer thinking about HCI, and one that explicitly encourages a more diverse and mutually constructive approach.

Although HCI can claim to be interested in almost any human activity because any activity can be mediated by computer (whether desktop, implant or immersive) what distinguishes HCI from the mere conjunction of human and computer is a concern for usability. However, whereas traditional sciences have a paradigm in which to scope legitimate objects and procedures of study, HCI has no such paradigm encompassing all of its subdisciplines. Ethics as a substitute underpinning for HCI activities has much to offer in terms of clarifying, contextualising and motivating HCI work.

Moreover, ethical complementarity is more than just a cute, ameliorative take on disagreements in HCI. Ethics as an explicit discipline has the benefit of at least two thousand years of effort. Ethics has traditionally been concerned with social and

individual good — without the added distractions and complexity that computer mediation brings to human interaction. Yet after thousands of years, there is no ready consensus or preference for any particular ethics. Instead an ethicist would first be clear which moral framework they were adopting (e.g., utilitarian, Marxist, virtue, discourse ...) *or* they would adopt the meta-ethical position.

Through interpreting and adapting ethics in HCI, it is hoped that HCI practitioners can find a suitable foundation in which to continue to grow knowledge and to do *good* work.

Acknowledgements

Harold Thimbleby is a Royal Society-Wolfson Research Merit Award Holder, and acknowledges their support. The authors are also grateful for very constructive comments from Jason Brotherton and Jeremy Gow.

REFERENCES

[1] ACM, ACM Code of Ethics and Professional Conduct, `http://www.acm.org/constitution/code.html`, 1997, updated 2003. See also D. Gotterbarn, K. Miller, S. Rogerson, "Software Engineering Code of Ethics is Approved," *Communications of the ACM*, **42**(10):102–107, 1999

[2] D. Alexander, *Rebuilding the Matrix*, Lion Publishing, 2001

[3] Aristotle, *Categories*, ∼300BC

[4] J. S. Brown & P. Duguid, *The Social Life of Information*, Harvard Business School Press, 2000

[5] O. K. Burmeister, "HCI Professionalism: Ethical Concerns in Usability Engineering," 2nd Australian Institute of Computer Ethics Conference, in J. Weckert, ed, Conferences in Research and Practice in Information Technology, **1**:11– 17, 2000

[6] S. Card, T. Moran & A. Newell, *The Psychology of Human-Computer Interaction*, Lawrence Erlbaum Associates, 1983

[7] J. M. Carroll & M. B. Rosson, *Usability Engineering*, Morgan Kauffmann, 2002

[8] M. Chalmers *et al.*, "HCI Theory is Like the Public Library," October, 2002 `www.chiplace.org/content/Hot_Topics`

[9] C. M. Christensen, *The Innovator's Dilemma*, Harvard Business School Press, 1997

[10] P. Duquenoy & H. Thimbleby, "Justice and Design," IFIP Conference on Human-Computer Interaction, in *Human Computer Interaction — Interact'99*, M. A. Sasse & C. Johnson (eds), pp281–286, 1999

[11] A. Einstein, *Relativity: The Special and the General Theory*, Routledge, 1994

[12] B. J. Fogg, *Persuasive Technology*, Morgan Kaufmann, 2003

[13] S. Fuller, *Kuhn vs Popper*, Icon Books, 2003

[14] M. R. Garey & D. S. Johnson, *Computers and Intractability*, W. H. Freeman & Co., 1979

[15] K. Höök, "Did you get rejected from CHI 2003?" November, 2002 `www.chiplace.org/content/Hot_Topics`

[16] B. Hoose, ed, *Christian Ethics*, Cassell, 1998

[17] T. Kuhn, *The Structure of Scientific Revolution*, 2nd ed, University of Chicago Press, 1970

[18] H. Küng, tr. J. Bowden, *Global Responsibility: In Search of a New World Ethic*, SCM Press, 1991

[19] B. Lawson, *How Designers Think*, 3rd edn, Architectural Press, 1997

[20] A. Lightman & O. Gingerich, "When do Anomalies Begin?" *Science*, **255**(5045):690–695, 1991

[21] J. Long, "A Discipline for Research Needs in Cognitive Ergonomics," *Theoretical Issues in Ergonomics Science*, **2**(3):289–308, 2001.

[22] Luke, 10:27–37

[23] A. Monk, "Fun, Communication and Dependability: Extending the Concept of Usability," in X. Faulkner, J. Finlay & F. Détienne (eds), *People and Computers*, **XVI**:3–15, 2002

[24] G. E. Moore, "Cramming More Components onto Integrated Circuits," *Electronics*, **38**(8):, 1965

[25] D. Norman, *The Design of Everyday Things*, MIT Press, 1998

[26] D. Norman, "Affordances, Conventions and Design," *Interactions*, **6**(3):38–43, 1999

[27] W. G. Perry, Jr., *Forms of Ethical and Intellectual Development in the College Years*, Josey Bass, 1999

[28] R. W. Picard, *Affective Computing*, MIT Press, 2000

[29] S. Pheasant, *Ergonomic, Work and Health*, Palgrave, 1991

[30] C. S. Pierce, cited in D. E. Knuth, *Things a Computer Scientist Rarely Talks About*, CSLI Lecture Notes **136**, CSLI Publications, 2001

[31] T. Pratchett, I. Stewart & J. Cohen, *The Science of Discworld*, Ebury Press, 2000

[32] J. Preece, H. Sharp & Y. Rogers, *Interaction Design*, John Wiley and Sons, 2002

[33] J. Raskin, *The Humane Interface*, Addison-Wesley, 2000

[34] J. Rawls, *A Theory of Justice*, Oxford University Press, 1993

[35] F. B. Schneider & M. Rodd, eds, *International Review of UK Research in Computer Science*, Institution of Electrical Engineers, 2001

[36] H. Simon, *The Sciences of the Artificial*, 3rd edn, MIT Press, 1996

[37] L. Suchman, *Plans and Situated Actions*, Cambridge University Press, 1987

[38] H. Thimbleby, "Design for a Fax," *Personal Technologies*, **1**:101–117, 1997

[39] H. Thimbleby, "Calculators are Needlessly Bad," *International Journal of Human-Computer Studies*, **52**(6):1031–1069, 2000

[40] H. Thimbleby, "Healthy HCI?" *ACM SIGCHI Bulletin*, p8, November/December, 2002

[41] H. Thimbleby, "Explaining Code for Publication," *Software—Practice & Experience*, in press.

[42] H. Thimbleby, P. Cairns & M. Jones, "Usability Analysis with Markov Models," *ACM Transactions on Computer Human Interaction*, **48**(2):99–132, 2001

[43] H. Thimbleby & M. Jones, "Obituary for a fax," *Personal Technologies*, **6**(2):151–152, 2002

[44] P. Thomas & R. Macredie, "Introduction to the New Usability," *ACM Transactions on Computer Human Interaction*, **9**(2):69–73, 2002

[45] P. Wegner, 'Why Interaction is More Powerful than Algorithms,' *Comm. of the ACM*, **40**(5):80–91, 1997

[46] I. H. Witten & D. Bainbridge, *How to Build a Digital Library*, Morgan Kaufmann, 2002

[47] R. M. Young, "Surrogates and Mappings: Two Kinds of Conceptual Models for Interactive Devices," In D. Gentner & A. L. Stevens (eds), *Mental Models*, 35–52, Hillsdale, NJ: Lawrence Erlbaum, 1983