Ethical Challenges in Medicine and HCI

George Buchanan
City University London
Northampton Square
George.buchanan.1@city.ac.uk

Second Author Name
Affiliation
Address
e-mail address
Optional phone number

ABSTRACT
Ethics are intended to reduce the potential for harm as a consequence of research activity. However, there are competing understandings of what constitutes harm, particularly in interdisciplinary research. We review three areas in which we have engaged in research that bridges medicine and HCI and identify differing ethical approaches that lead to outcomes and compromises that may themselves lead to harm by inaction. We suggest some paths for further reflection that may help build processes that better support the aims of ethical research.

Author Keywords
Ethics, medical practise, interaction design

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Health is an area of growing interest in human-computer interaction research. However, investigating the domain brings into direct contact very different approaches to ethics. In medicine, the issue of physical hazard is critical to the ethics process, and there are concerns about patient confidentiality founded on the context of many procedures and conditions being potentially embarrassing or sensitive. In HCI, while there can be physical hazard, potential embarrassment and sensitivity of topic, this is not such a consistent theme.

HCI faculties have regularly encountered changes in ethical procedures, and incorporated concerns from (e.g.) medicine, health, psychology and the social sciences. While such changes have been seen to be – and occasionally have been – applied uncritically, there have been benefits in considering user's concerns and needs more consistently than was the case in (say) the 1980s.

However, medicine and health research in HCI, often subject to approval through medical faculties, has in many cases shown particularly extreme cases of both the mis-application of the clinical paradigm, but also a lack in the ethical approval cases in medicine and health.

There are two major paradigms to ethical action in research: pre-emptive and in-context. Many subjects combine these, but almost all emphasise pre-emptive ethics over in-context, as these are seen as more planned and principled. In other areas, where observation and research often uncovers unexpected ethical challenges, in-context ethical action, with or without a formal overview process, is more commonplace.

The CHI+MED project is a large (AU$11m) human-computer interaction project that involves four different universities. Different members of the team have faced individual challenges with obtaining and negotiating ethical approval, but also in encountering unexpected situations and ethical considerations.

INVESTIGATING CALCULATION ERRORS
The CHI+MED project was motivated by research conducted in the medical domain that reported that errors in controlling and programming clinical devices were a major contribution to avoidable harm in patients. One specific motivating case occurred in a hospital in Alberta, Canada, where a patient was given a 96 hour (4 day) dose in a 4 hour period[1]. While the medicine was strong, but not harmful, in the longer period of administration, the concentration delivered was fatal.

A key aim from the start of the project was to understand how human factors and interaction designs could be contributing to such errors. With a better understanding, our hope was to ensure that clinicians and their patients could be better protected against human fallibility, even where great care and attention is already being taken.

The very large number of device programming tasks undertaken each day – several million in the U.K. alone – naturally means that even a very low rate of error will, in aggregate, lead to a total figure that would be a matter of concern. Observation in hospitals would itself have ethical concerns: it might lead itself to increased rates of error, by distracting staff, and furthermore the low rate of error would mean a substantial cost for very little data gained. Combined, these factors led us to discount such an 'in vivo' approach as unprincipled.

Retrospective error reporting already is undertaken (e.g. Johnston, 2006), and thus a 'critical incident' approach to investigating errors could also be used. A small sample of these reports was made available to the project, with the

OzCHI '15 , December 07 - 10 2015, Melbourne, VIC, Australia
Copyright © 2015 ACM 978-1-4503-3673-4/15/12... $15.00
http://dx.doi.org/xx.xxxx/xxxxxxx.xxxxxxx.

[1] http://www.cbc.ca/news/technology/fatal-chemo-overdose-prompts-alberta-reforms-1.650904

generous assistance of clinicians. However, there is always a significant risk in such cases of post-rationalisation by parties to the case, and furthermore the data available to us lacked the materials necessary for an interaction-focussed investigation. As interaction issues have not been a routine matter of concern, there is a lack of expertise amongst investigators, who have seldom received training or education on HCI and related matters.

There was a further suite of ethical issues regarding any potential involvement with ongoing investigations, which precluded that approach. Our involvement might compromise or impede current good investigative practice, and we could potentially reveal information that implicated one or other party to the incident – be that members of clinical staff, providers of devices, support staff or otherwise.

One major practical problem in investigating the potential calculation errors made in clinical practice is the underlying assumptions of the broader clinical professions regarding error and intention. Training is a critical element of good practice in medicine. Compliance to process is understood - even, perhaps, presumed - to either eliminate, or to reduce to a practical non-entity, human error. Deviation from process is seen as indicating a lack of professional care. However, while lack of attention to process certainly risks patient harm, practical stresses of timeliness, workload and cognitive overload appear to be overlooked. While we know, from psychology, that cognitive duress naturally and unavoidably contributes to human error, this view is not built-in to medical investigations of error.

A comparison with aviation is helpful: there, while training is certainly seen as critical, and breach of process has contributed to fatal incidents, such as the Air France disaster, it is also recognised that good practice may involve deviation from routine, and that cognitive factors are powerful agents that can overwhelm even diligent professionals.

However, drawing on the long-established practice of using artificial contexts, a natural approach would be to investigate error rates in non-clinical settings. Ideally, we could use tasks, contexts and distractors that are analogous to, but do not carry the immediate risks of, medical work.

The calculation of dosing is of course a critical skill in nursing. The ability of nursing staff to undertake accurate calculations has been researched, most often using trainees. Researchers in the education of nurses at Swansea University discovered, (Harvey et al 2009), that only 19% of a cohort of trainee nurses achieved a score that indicated proficiency in basic mathematics. While training would undoubtedly increase the proportion achieving that result, it is nonetheless a matter of concern that so few nurses enter the profession with proficiency in basic mathematics.

This information fell neatly into the avenue of research that we pursued within CHI+MED to reduce dosage errors, including calculations made by nursing staff.

Taking an experimental approach, undertaking calculations in a non-medical environment, we ideally wanted to perform experiments and other investigations that included the target user group (i.e. nursing staff). The calculations performed would not be administered doses, nor part of any clinical procedure. The calculations would therefore have no direct clinical repercussions.

However, there was a reticence from a number of institutions and parties to this area of research. There were concerns that if an individual performed poorly in a test, or a cohort did, that that knowledge would necessarily have ethical implications for practise: i.e. that patients may in fact be being put at risk by the low proficiency of individuals, but also that individuals might be placed in a position of embarrassment if they made errors. Thus, there were practical barriers to achieving a goal of improving patient safety.

In this case we could readily recruit participants of equivalent mathematical ability to the 'typical' nurse, but we suspect that in practice, nursing staff have developed work-arounds and defensive measures to reduce the likelihood of error, and increase patient safety. Understanding those strategjes, and indeed reinforcing them where possible could, arguably, enhance patient safety and the professional security of clinical staff.

What we did do in the final event was to recruit lay-people whose mathematical ability was analogous to that of trainee nurses. Some of the potential methods we experiment with for reducing error proved effective, but the exercise also revealed that there was a weak sense of what the 'right' answer should be, meaning that several errors were made in the cohort, even with the best designs.

SELF-MANAGEMENT

As medicine moves from being delivered by professionals to increasing degrees of self-monitoring and self-management, new issues emerge both practically and ethically. One major strand of research in CHI+MED has been the design of novel technologies to support patients oversee their own health. Chronic conditions are the dominant element of the move to self-care, with type 2 diabetes in particular being a major concern for national health authorities.

Returning briefly to the issue of dosage, where error by clinicians is possible, in relatively controlled conditions. The medical professions are well aware that there is a significant question as to whether patients can be expected to achieve even similar levels of performance. Undertaking experiments with patients would be possible, but there are few direct comparisons. The dosage regime followed with patients is altogether different. Patient-administered doses are often given in pre-set dosage sizes, with known effects. Where dosage varies across a wider range, in the vast majority of cases this is guided– or ideally controlled–by dose within pre-calculated safe ranges. The discretion of the patient is in using a higher or lower dose within that range (this is the norm for insulin, for example). Errors can still accumulate – e.g. over-dosing can occur when previous treatments are

forgotten. However, they are not calculation errors, but rather errors in tracking previous treatment. Furthermore, the safety implications of many of these drugs are more limited, as the tolerances between the normal dose and a potentially harmful dose is usually much greater than in some of the treatments clinicians have to administer.

Turning to a broader perspective, self-care holds a complex and ambivalent place within existing frameworks. Drugs and devices that directly interact with the patient are, quite understandably, exposed to high levels of inspection and ethical oversight. However, it is clear from our interaction with device manufacturers, procurement officers and medical staff, that usability is not a factor that is typically included in the procurement process, and opportunity for user error can be introduced in subtle ways.

In one case, a new form of diabetes tool required users to 'pull' the device, rather than 'push' it, even though the device *afforded* no clue as to its different operation. For a number of reasons, this contextual usability problem was not identified before a strategic commitment was made to move entirely to the new tool for all patients. Reproducing this sort of situation in a usability study would produce all sorts of ethical dilemmas – from potential (minor) injury to ineffective or inaccurate self-treatment. However, the deployment in practice, without usability analysis, reproduced the same risk in actual practice, and it took some time for the problems patients were encountering to emerge, due to reticence to admitting difficulties, and the fact that many patients had existing stocks of old materials that they exhausted before adopting the new devices.

An underlying complication in both ethics and practice – as revealed in this example – is that human factors, which are potentially even more significant in occasional or partially trained use, form a facet of overall safety that the medical profession associates with personal error. Error is seen as being primarily eliminated via compliance to procedure. Consumer-facing devices in other areas - such as mobile interaction - increasingly support permissive interaction, where the enforcement of single process pathways is seen as undesirable. We simply do not have an adequate evidence-base upon which any ethical process can be built. We do not know the rate of human error that pervades consumer-facing devices, we do not know the potential gravity of the consequences of human error, and we do not possess a rigorous framework for assessing individual culpability or responsibility. This raises a rather alarming double-barrelled question: how can we design effective patient support, without either presuming a very high (and undefined) level of personal responsibility, or that any attempt to create self-management should be avoided (with the con-commensurate risks to the well-being of the populace as a whole).

EMERGING CHALLENGES

Beyond these existing issues – some of which have certainly been rehearsed before, there is a new category within which there is an ethical vacuum. 'Health' or 'wellbeing' apps of various forms are exempt from mandatory compliance procedures in many legislatures. Conversely, some universities can perceive any health-related investigation as inherently requiring the same ethical standards as the research of drugs and devices.

In the case of chronic conditions, we have engaged in a number of ways with patients, within normal HCI ethics processes. While obtaining ethical consent was not particularly complex, there appeared to us to be issues of contextual ethics that the formal process overlooked, and where neither the clinical nor traditional HCI or computer science approach to ethics appeared to function effectively.

One research activity included delivering an intervention that would assist users in self-management of their diabetes (Owen et al 2015). One concern that we had, from discussions with clinical practitioners within the broader project team, and with other researchers, was the issue of legacy support, following the research. Nowhere in the ethical process was the issue of how to handle cessation of the study. Two participants (of twelve) asked to retain the app we had developed after the study, and incidental and unplanned reporting revealed that two other participants used it afterwards without direct discussion with us. At no stage in the ethics process were we asked to consider this issue, yet we deemed it important to address from the beginning of the project, and we had a planned exit strategy from the outset (including a period of limited support if requested).

CONCLUSIONS

While it is easy to identify shortcomings in both clinical and HCI ethics procedures, these are in many ways an outcome of a quickly moving context in which formal procedures naturally lag behind new and emerging problems. HCI is engaging with areas in which there are new risks, and yet it also is aware of problems and difficulties that medical ethics have construed in ways that seem from an interaction and cognitive perspective to be ill-informed. Gaps also exist in which both approaches are inadequate, and an ongoing dialogue is required if effective and ethical research is to be done. However, it must be borne in mind that research has demonstrated that current techniques and systems are themselves causing harm, and inaction is itself almost certainly unethical.

ACKNOWLEDGMENTS

Many thanks for the support of the Engineering and Physical Sciences Research Council for their support.

REFERENCES

BBC News (2011) http://www.bbc.co.uk/news/uk-wales-south-east-wales-13625168

Harvey S., Murphy F., Lake R., Jenkins L., Cavanna A., Tait M. Diagnosing the problem: Using a tool to identify pre-registration nursing students' mathematical ability, Nurse Education in Practice 10.3. June 2009, 119-25.

Johnston, A.M., Unintended overexposure of patient Lisa Norris during radiotherapy treatment at the Beatson Oncology Centre, Glasgow in January 2006. Scottish Executive, 2006.

Oladimeiji, P., Pearson, J., Owen, T., Buchanan, G., Towards improving error detection: an example in the design of calculators, *in press*.

Owen, T., Pearson, J., Thimbleby, H., and Buchanan, G.. ConCap: Designing to Empower Individual Reflection on Chronic Conditions using Mobile Apps. *Procs. Intl. Conf. on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '15). ACM, (2015) 105-114.

Thimbleby, H.. Permissive user interfaces. *International Journal of Human-Computer Studies*, 54.3 (2001), 333-350.

The columns on the last page should be of approximately equal length.
Remove these two lines from your final version.