

The Pain of Crowds: Considering Wider Ethical Implications in Conference Planning and Review

Dana McKay

University of Melbourne
Parkville 3010 Australia
dmckay@swin.edu.au

George Buchanan

Centre for HCI Design, City University
Northampton Square, London, UK
George.buchanan.1@city.ac.uk

ABSTRACT

With the increasing availability of data, an increasing social awareness of the impact of individual trauma, and the increasing penetration of HCI into sensitive areas, ethical questions in HCI are more challenging and prevalent than ever. These issues certainly affect the practitioners of HCI in sensitive settings, but also have the potential to affect reviewers and even conference attendees. Moreover, with the challenges raised by big data, it is more important than ever that reviewers are ethically vigilant. This position paper addresses some of the challenges associated with reviewing ethically fraught work, with a view to positioning the responsibility of ethics not just with committees and individual researchers, but with the HCI community as a whole.

Content warnings

Sexual assault, sexism, military activity, classification of human beings, death.

Author Keywords

Add your own keywords here

ACM Classification Keywords

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

INTRODUCTION

It is abundantly clear that ethics in the field of HCI is becoming a more complex issue to address. More and more people are users of technology, and some of these users are in some way a vulnerable population (for example children or people with disabilities). The contexts in which technology is used are also expanding, into areas where HCI has previously not ventured—for example health and medicine and even the bedroom (Bruckman 2014).

At the same time, the amount of publically available material that could be used as data for studies of human behaviour or technology has increased substantially: we are in a ‘data deluge’ (Borgman et al. 2007). The use of publically available data is often not covered or restricted by the traditional method of ensuring ethical research: the institutional review board or ethics committee (Bos et al. 2009).

Ethical Encounters: HCI Research in Sensitive and Complex Settings, an OZCHI workshop, Dec 7 2015, Melbourne, Australia.

Copyright in material reproduced here remains with the author(s), who have granted the workshop organisers a perpetual, non-exclusive license to distribute it via the workshop website (<https://ethicalencountershci.wordpress.com/>). For any further use please contact the author(s).

Finally, we are increasingly aware that individuals—participants, fellow researchers, reviewers and conference attendees—may carry with them the effects of trauma, and these traumas may affect their experience of being involved with research. Such effects are known to occur in workplaces in general (United States Congress 2003).

Both authors of this paper have a significant though not well publicised interest in social justice; we are both versed in literature on diversity and discrimination in a wide range of forms. This gives us a unique perspective on the field, and we have found we are often arguing for something we feel is an ethical ‘greater good’ in review and conference organisation, where such has not been considered by paper authors, fellow reviewers and conference organisers.

This position paper draws on our experience of arguing for ethics and presents a trio of case studies where conference organisation processes threw up ethical issues that fall outside of the traditional ethics process, then reflects on how we as a community might address these.

CASE STUDY ONE: PUBLIC WORK ON SEXUAL ASSAULT

The ethical issue in this case was raised by a paper—that was eventually published—on the collection of publically available materials on the accusations surrounding Jimmy Savile (Llewellyn et al. 2014). The paper itself noted that there are ethico-legal issues about creating an archive of such material, because Savile died before he could face trial for his alleged crimes, but the paper presented further ethical issues than just those. One of us was asked to review this submission for the conference. Additionally, due to the sensitive nature of the paper the review was requested by email—rather than through the reviewing system—so that not even the usual handful of conference organisers and co-reviewers who would normally know the identity of reviewers would know who reviewed this particular paper. This anonymous approach was used so that should any of the reviewers be personally affected by sexual assault, they could reveal this in their review process. Ideally under such circumstances, program chairs would ensure that the paper was reviewed by survivors to ensure it was not likely to cause them harm (Charlton 1998). There are ethical challenges, however, both in sourcing such reviewers among the subject matter experts on a program committee (one cannot ask survivors to self-identify) and in asking a small number of survivors to represent the views of a group.

The greatest challenge in this review process, though, was that the paper was presented as a poster submission, requiring public exposure of a sensitive issue at an

unexpected venue. The number of sexual assault survivors in the general population is as many as one in four women and one in seven men (Wikipedia 2015); at a conference with between two and three hundred attendees it is a statistical certainty that some will be survivors. Of these, a significant number may experience PTSD and thus be susceptible to being triggered (RAINN 2009); certainly the conference in question has not presented material of this nature in the past and attendees would have no reason to expect to encounter it. In such cases, trigger warnings or content notes can be useful (Chemaly 2014). Some may argue that sensitivity to survivors of sexual assault should not trump academic freedom (Lukianoff et al. 2015), but it seemed to us that a balance could be struck by converting the submission to a short paper; the submitted review put this case quite strongly. The remaining two reviewers considered the merits of the paper, but not these broader implications for the community. While the change to a short paper did not eventuate, the program chairs handled the matter sensitively, and measures were put in place to both allow the publication of what was, academically, an important paper and minimise the risk to survivors of sexual assault. These measures included a content warning given during the minute madness session of the conference, no explicit material in the poster itself, and obscuring the online location of the data repository on which the poster was based.

CASE STUDY TWO: AUTOMATIC CLASSIFICATION OF HUMAN BEINGS

As this case describes a paper that was not accepted for publication, it will describe the situation involved in generalist terms. The paper in question was about the use of data to automatically classify human beings on the basis of their names according to what is referred to in law as a ‘protected characteristic’ (Equality and Human Rights Commission 2015). One of us has an interest in the use of ‘big data’ and volunteered to review the paper on that basis.

The ultimate goal of this paper was an ethical good: to determine minority representation within an academic field that has a number of significant biases; however the approach used by the authors was a data-driven machine learning approach. Given that the paper never involved the use of live human subjects, but only publically available data, it would be within the remit of any typical ethics process (Hearnshaw 2004; Bos et al. 2009), and therefore the authors had clearly not sought ethical advice. Nonetheless, the paper both named individuals and their assigned classifications, and drew conclusions about the field it studied on the basis of these classifications. There were ethical problems within the methodology of the paper as well, in that the machine learning algorithm drew on Wikipedia (itself hardly a representative source (Callahan et al. 2011; Greenstein et al. 2012)) as a corpus. Furthermore the method presented in the paper was not compared to the generally accepted method, which assesses diversity using a well-accepted substitute for the protected characteristic in question.

Two of the reviewers of the paper were data specialists with no particular interest or experience in human computer interaction; the third was one of us. The data specialists rightly assessed the paper as providing an interesting contribution to the field; they were quite in favour of its inclusion in the conference. Stacked against this contribution to data processing, it is ethically challenging to classify people, especially if some of them are named in this paper. This difficulty was intensified by the fact that (according to the accuracy statistics presented in the paper) the algorithm described was fairly likely to misclassify people. Misclassification is well known to be at least hurtful, if not actively harmful (Stepanikova 2010; Maza 2015). The chance of harm does not, in our opinion, mean that such work should not proceed—just that care should be taken to ensure that it is conducted as ethically as possible. In this instance, such care was not taken: the researchers involved named individuals, did not use a representative corpus, did not compare their algorithm to the nearest available (and less ethically fraught) alternative, and had evidently not considered the ethical implications of this work. On the strength of the ethical objections to this work, it was ultimately rejected from the conference to which it was submitted.

CASE STUDY THREE: THE ASSANGE CASE

The choice of keynotes at conferences is often a decision made after extensive discussion. Ideal keynotes should make a contribution to the conference as a whole by stirring up new ideas and generating further conversations outside the keynote session.

The British Computer Society has organised one of the longest-running conferences in HCI, which in 2014 was in its 28th year. In 2015, the organisers of the conference decided to invite Julian Assange—a well-known and, clearly, controversial speaker. With the inclusion of a sponsor organisation, which has a formal policy on professional ethics in information technology, this particular invitation raised significant problems for the sponsor. The BCS code of conduct is publicly available (British Computer Society) and identifies the key principles that its members are expected to uphold.

One section addresses the issues of the handling of confidential data. Julian Assange is well-known for his involvement in WikiLeaks, and in particular his argument for the policy of not making any redactions. This has led to claims that third parties might come to harm (e.g. Afghani interpreters working for western governments). Nick Cohen of the Guardian newspaper, who spearheaded the release of redacted documents, revealed Assange’s response to questions about possible risks to this group, in 2011 ‘if they get killed, they’ve got it coming to them. They deserve it’ (Cohen 2011).

This is clearly not consistent with the BCS policy of protecting third parties, which applies even in cases where the public interest argument means that the normal duties of privacy of data are otherwise suspended. The safety of every member of society is also a more general underpinning principle.

A separate issue was the charges on which Sweden sought to investigate Assange. While it is important to uphold the principle of not presuming guilt, the circumstances raised two practical issues: first, Assange could not address the conference in person; second, his presence—virtual or otherwise—could serve to exclude a number of groups. In terms of the keynote, the ethos of the conference has been to have an in-person keynote, as a keynote’s contribution should not be limited to the hour in which they make their address.

However, the involvement of a figure who faced investigation regarding crimes of a sexual nature, and who had fled those accusations, might serve to divide and exclude part of the community. Again, the BCS code of conduct does not simply address representation of under-represented and vulnerable groups, but to ‘actively promote’ the inclusion of all people. Furthermore, the profile of the conference leadership set additional context. Whereas in previous years, women researchers were well represented in both program- and general-chair roles, in 2015 the team was comprised entirely of men, and at the time of the announcement of Assange as keynote, the other keynote was also a man.

The positive arguments for Assange, in addressing major societal issues could not be simply set aside, and there were also concerns that refusing his participation would be a suppression of academic freedom. Thus, the circumstances pitched one public concern—holding governments to account—against the inclusion of a number of groups, including women and survivors of sexual assault. This naturally led to a heated discussion of how to proceed.

The BCS, as sponsor, was not willing to proceed with a conference that failed to uphold two key principles of its code of conduct (which are part of its charity charter, thus the breach of them would threaten its charitable status).

The local organisers determined to proceed with Assange as keynote, with the BCS withdrawing sponsorship, and several research groups, particularly from the north of Europe, who regularly sent members, refusing to participate. The organised event did occur in July 2015, with Julian Assange as one of the keynotes, but not as a part of the BCS HCI series; that series continues with a conference in 2016.

During the discussion among the decision-makers involved, some argued that those who found Assange an uncomfortable proposition were guilty of prejudging the investigation, regardless of any further concern with the data protection issue. That group repeatedly declined to address how they would support anyone with personal history that was related to the charges on which Assange was being investigated, and some suggested that any negative reaction was a personal problem, not a matter for the community as a whole to address.

Conversely, others argued that the principle of inclusion of a wider group was more important than the privilege of one person to present a keynote, and that the topic being addressed could be done by other means. Given the proportion of individuals who are survivors of sexual

assault, a conference of 300 or more participants could certainly expect a sizable number to be affected by the issue.

LEARNINGS AND CONCLUSIONS

The experiences presented in this paper have all happened in the course of a year; while it is unusual that so many ethical challenges should appear in such a short time, it has created a lens for the authors to think about broader ethics in HCI.

The first clear learning is that every time we—as individual researchers, and as a discipline—study human beings and write about them, we incur the possibility that we will affect not just our participants, but those who review, hear, or read our work. When we are dealing with sensitive matters, such as sexual assault, we should consider all of these groups in the way we present our work.

The second learning is that the motto ‘nothing about us without us’ (Charlton 1998) should apply not just in policy social justice work, but in academic circles as well. This phrase is social justice shorthand for not making decisions that affect marginalised groups or individuals without the input of those groups or individuals; such behaviour is paternalistic and may well cause more problems than it solves. Observing this principle may mean empowering our participants as co-authors, or ensuring that underrepresented groups are included in review panels and organisation teams, and consulted as appropriate.

The third learning is not new: we cannot rely on an ethics process to ensure research is conducted ethically (Bruckman 2014), nor can we rely on ethics policies of national bodies to ensure ethical behaviour. The implication of this, though, is that we as researchers who care about ethics must be vigilant, and we should speak up wherever in the publication process we encounter ethical concerns. In all three cases presented here, one or two people speaking up effected change in a way that supported disadvantaged groups.

The final learning, on which we shall conclude, is that this work is hard. It can involve feeling exposed as a member of a minority group, or indeed revealing hidden disadvantages (such as disability or personal history) that one would rather not reveal. We feel, though, that it is also worthwhile; without people who do this work ethically problematic work will be published, and HCI culture will not change. If we move away from the ‘ethics process’ to actually caring about ethics throughout the research process from planning to publication, we will do better, more humane work and be a better, more inclusive discipline.

REFERENCES

- Borgman, C., Wallis, J. and Enyedy, N. Little science confronts the data deluge: habitat ecology, embedded sensor networks, and digital libraries. *IJDL* 7, 1, (2007) 17-30.
- Bos, N., Karahalios, K., Musgrove-Chavez, M., Poole, E. S., Thomas, J. C. and Yardi, S. Research ethics in the

- facebook era: privacy, anonymity, and oversight. Proc. CHI 09 Extended Abstracts, ACM. (2009), 2767-2770.
- British Computer Society. Trustee Board Regulations Schedule 3 Code of Conduct for BCS Members. British Computer Society.
- Bruckman, A. (2014). Research Ethics and HCI. Ways of Knowing in HCI. J. S. Olson and W. A. Kellogg, Springer New York: 449-468.
- Callahan, E. S. and Herring, S. C. Cultural bias in Wikipedia content on famous persons. JASIST 62, 10, (2011) 1899-1915.
- Charlton, J. L. Nothing About Us Without Us: Disability Oppression and Empowerment. Los Angeles, CA, California University Press (1998).
- Chemaly, S. What's Really Important About 'Trigger Warnings'. Huff Post College. Huffington Post (2014).
- Cohen, N. The Treachery of Julian Assange, The Guardian Comment is Free (2011). Available at: <http://www.theguardian.com/commentisfree/2011/sep/18/julian-assange-wikileaks-nick-cohen>
- Equality and Human Rights Commission. (2015). "Protected Characteristics." Retrieved 11 September 2015, from <http://www.equalityhumanrights.com/private-and-public-sector-guidance/guidance-all/protected-characteristics>.
- Greenstein, S. and Zhu, F. Is Wikipedia Biased? The American Economic Review 102, 3, (2012) 343-348.
- Hearnshaw, H. Comparison of requirements of research ethics committees in 11 European countries for a non-invasive interventional study. BMJ 328, 7432, (2004) 140-141.
- Llewellyn, C., Ruus, L., Burnett, R., Kirkwood, S., Smith, M. and von-Jungenfeld, R. Building a dataset of sensitive information. Proc. DL 14, IEEE Press. (2014), 493-494.
- Lukianoff, G. and Haidt, J. The Coddling of the American Mind, The Atlantic (2015). Available at: <http://www.theatlantic.com/magazine/archive/2015/09/the-coddling-of-the-american-mind/399356/>
- Maza, C. Journalism Experts Condemn NOLA.com's Misgendering Of Transgender Shooting Victim, Media Matters for America Blog (2015). Available at: <http://mediamatters.org/blog/2015/02/13/journalism-experts-condemn-nolacoms-misgenderin/202528>
- RAINN. (2009). "PTSD (Post Traumatic Stress Disorder)." Retrieved 11 September 2011.
- Stepanikova, I. Applying a status perspective to racial/ethnic misclassification: Implications for health. Adv. Group Proc. 27, (2010) 159-183.
- United States Congress. Congressional Record October 24-November 4. Congress. Government of the United States of America (2003).
- Wikipedia. (2015, 30 August 2011). "Sexual Assault." Retrieved 11 September 2011.