
Studying studying – academic life as the subject of commercial research

Alan Dix

Talis
Birmingham, B1 3HN, UK
and
HCI Centre
School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK

alan@hcibook.com

<http://alandix.com/academic/papers/chi2017-remix-studying/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Request permissions from Permissions@acm.org.
CHI'17 Extended Abstracts, May 06-11, 2017, Denver, CO, USA
© 2017 ACM. ISBN 978-1-4503-4656-6/17/05...\$15.00
DOI: <http://dx.doi.org/10.1145/3027063.3053685>

Abstract

This paper offers a personal view on the interplay between the multiple roles of academic and commercial researcher and classroom teacher. A series of vignettes expose a variety of lessons. Some of the vignettes are about the way that being an educator of human-computer interaction has informed my practice as an academic researcher about HCI. Some are about the different ethical expectations when using educational practice as part of commercial research as opposed to academic research. We also see how the inside knowledge of being a practitioner researcher offers unique insights into rich data and how product orientation can expose gaps in research, surprisingly not so different from the way that teaching students does: with neither developers nor students can you hide fuzzy thinking behind long words. While there are conflicts that need to be carefully managed, both academic and commercial research have powerful synergies with educational practice.

Author Keywords

Educational technology, research ethics, learning analytics, academic life.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; K.3.1 Computer Uses in Education

Sub-seat woofers today

While writing this paper, I tried a web search for sub-seat woofers, but still cannot find anything written about them. I assume the term is gamer community insider language.

However, there are a number of commercial products making use of the effect including ButtKicker (sic). Another company, SubPac, are building similar low-frequency transducers into a variety of form factors including rucksacks as well as gamer chairs. I am sure that there are more examples in the adult entertainment industry.

Interestingly, while there was extensive prior work on audio and sound in the HCI community, in the late 1990s, when my students were teaching me about sub-seat woofers, research on tactile feedback was still relatively new.

Introduction – Juggling Roles

I have been involved in HCI research, for over thirty years, and as an educator, both textbook author and teacher for twenty-five years. These academic roles have sometimes been synergistic, sometimes in conflict. More recently, while retaining a part-time academic post, my principal employment has been as a senior researcher for Talis, a small company that creates teaching and learning software for higher education. I have also taken a strong practical interest in the application of technology for communities at the edge particularly on Tiree, where I run the biannual Tiree Tech Wave, and embodied in my 2013 one-thousand-mile walk around Wales.

This paper will focus principally on the interactions between the roles of commercial researcher and academic educator including pragmatic, ethical and theoretical factors. However, as this is the first CHI Remix, I will start broadly and discuss several vignettes of academic teaching–research interconnections before looking at the additional commercial issues.

Learning from Teaching – sub-seat woofers

The interplay between educator and researcher has often been interesting, with outcomes of research becoming topics taught, but also the other way round where lessons learnt in the classroom have informed research at a practical and theoretical level.

My favourite example was while marking an exam on visualisation and virtual reality in the late 1990s. The question asked students to list mechanisms to promote engagement and immersive experience, a topic that had been covered in one of the lectures. Several students started their list with 'sub-seat woofers', a

term, new to me, that I took to mean bass speakers placed beneath chairs to create vibration during car racing or air combat games. The students were experienced gamers and spoke from that experience.

Not only was this not in my course, but I realised that none of the mechanisms I had taught included sound. I then looked at the textbooks on multimedia available at the time and found that none devoted more than one or two pages out of several hundred to sound. Yet for these highly-experienced users sound topped their list.

Of course, there had been plenty of HCI research on sound including Gaver's early work on ecological acoustics [12] and Brewster's sonic enhancement or replacement of interface elements [3], but this had not found itself into the specialised textbooks of the time and certainly not into my teaching.

There are two lessons here. First being aware that unexpected answers may not be wrong answers ... even in the midst of exam marking. The other is that, albeit not always expert (despite the hype on digital natives [16, 17, 14]), students are avid users of technology and thus excellent to listen to. Of course, the students concerned also got top marks for the question!

From Teaching to Research

In the sub-seat woofer example, I was learning from the students as users. We'll now look at two examples where the act of teaching acted as a forcing function creating new research knowledge.

The first occurred during the teaching of the same module on visualisation and virtual reality. I had been teaching about 3D visualisation and how *static* 3D




Figure 1: Interactive Stacked Histograms or 'Dancing Histograms'. Normally only the overall sum and the bottom category are easy to see as timelines. In the interactive version, selecting an element in the histogram or key drops the histogram bars so that the category of the selected cell is on the baseline. This means that the trend of the selected category is easily readable. In this case, the slow increase in the yellow banana sales since 1993 is evident. (Image ©Alan Dix)

<http://www.meandeviation.com/dancing-histograms/>

visualisation was particularly problematic; you needed to be able to rotate or change perspective. I then made the statement that "*any visualisation can be improved by interaction*".

Having committed the academic's mortal sin of hyperbole, I needed to either justify or retract the statement. Happily, on the board was a drawing of a stacked histogram where earlier we had discussed the need to make choices during the design of static visualisations. So I was able, during the lecture, to re-design the stacked histogram by adding interaction to change the baseline making what we later called 'dancing histograms' (fig. 1). Geoff Ellis and I went on to create further examples of adding interaction to standard visualisations and express the general design principles for doing this in an AVI'98 paper [5].

The second example, occurred some years later early 2000s) teaching masters students about low-intention systems, such as automatic lights and doors that do useful things without direct interactions. I had talked about the issues before in research contexts [6] and effectively ended with a series of 'things we don't know'. While this was an acceptable end point for a research talk, it felt uncomfortable in a teaching context. This kind of system was, even at that time, being deployed, and it seemed insufficient to say, "wait 10 years until the research is ready".

As part of the course we had already made heavy use of detailed scenarios and traces as ways to connect different kinds of formal and informal descriptions. So we took one of the examples of low-intention interaction, car courtesy lights, and as a class exercise annotated traces of typical 'getting into the car'

behaviour as a way to help in practical design. This case study later went into the third edition of our HCI textbook [7], which then led to more detailed semi-formal analysis and step-by-step design guidance in online material for the book¹, and more recently into a chapter in the new Springer Handbook of Formal Methods in Human-Computer Interaction [11].

In some ways the teaching situation in both these examples acted in a similar way to a practical design situation. It was evident that a general statement needed to be transformed into concrete examples and advice. Furthermore, 'exposing open research questions', a good endpoint for research, needs to be turned into at least the beginning of a solution – and through this actually discovering more than would have felt possible from a pure researcher role.

Educational Experiments

Like most academics I have always experimented on my students, for example using web pages as lecture slides back in the mid 1990s. We adopt what we think are good methods or simply interesting ones, often with minimal prior empirical evidence of effectiveness.

At first this sounds ethically problematic; however, practice-based innovation is at the heart of many accepted methods for educational research and professional development [2]. Furthermore, teacher motivation and engagement is often one of the most critical aspects of student learning, so practices that add to the excitement or freshness of delivery are likely

¹ HCIBook online! Car courtesy lights - designing incidental interaction. <http://hcibook.com/e3/casestudy/car-lights/>

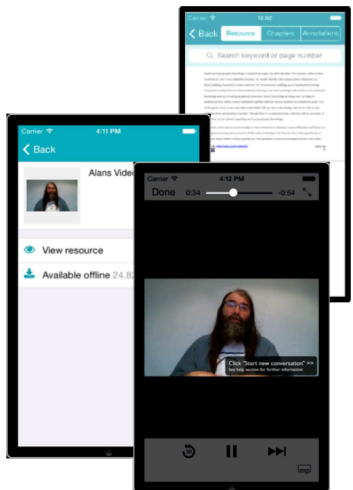


Figure 2: Talis Lighthouse pilot universal media player: (above) PDF in web interface ; (below) PDF and video in mobile app (image ©Talis)

to end up being effective teaching, almost irrespective of the intrinsic value of the methods chosen.

Having dual academic and industrial employment, there are clear synergies: from an academic point of view having 'inside' access to new products and from a commercial point of view being able to use new software 'in the wild'. However, with this come potential conflicts of interest and I have found myself taking a far higher degree of care than I would if I were using, say, experimental academic research software.

This was not a great issue when I used Talis Aspire Reading List (TARL) for a HCI mini-MOOC² in 2013 as this was well-established software on a free course and only my own reputation risked damage. However, using the Talis Lighthouse pilot (fig 2) with my university students was slightly more problematic.

The conditions that set my own mind at rest were:

- i. the fact that other academics were taking part in the same pilot programme;
- ii. openness to colleagues and students;
- iii. it was only used on small parts of modules, hence reducing the impact of any potential problems;
- iv. the modules were ones where I was not the primary academic, effectively operating under supervision.

Interestingly, other academics involved in the Lighthouse pilots, but who had no other connection to Talis, were not worried about (iii) and (iv), and I would

not have myself if it had been non-commercial research software, even though it is likely the latter would have been less stable and less carefully designed.

Publishing Results

The outcomes of this and on-going research have been widely published including implications for flipped class teaching, open-education resources, and learning analytics. However, it has also been interesting observing my own and reviewers' attitudes.

Just as with the execution of the research, the potential for conflicts of interest has influenced styles of writing, and I always try to make explicit both the real educational context for external validity, but also the commercial research interest. Similar conflicts arise in any educational research: educator and researcher share the same long-term goal of improving education, but may be at odds in the short term, especially in an age of metrics-based academic performance.

The same additional level of rigour seems to apply to reviewers. The most amusing case of this was a reviewer who thought the name of the systems being used were mentioned too often – which turned out to be a total of ten mentions in an eighteen page paper. Attempts to replace this with descriptive terms such as "the system being used" edged on the farcical. Clearly the reviewer would never have made the statement if this had been the name of a research system.

Insider Information

Although there are conflicts in being both researcher (whether academic or commercial) and practitioner, the experience of action or practitioner research is that usually there are also substantial gains.

² Now available at Interaction Design Foundation. <https://www.interaction-design.org/courses/human-computer-interaction>

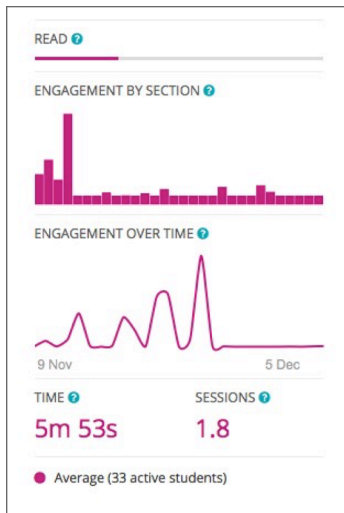


Figure 3: Summary analytics.

The histogram above shows the profile of pages from beginning to end with the average number of students who have looked at each page. Note that most student have only viewed the first few pages.

The line below is time graph of when students engaged with the resource.

(image ©Talis)

The Talis Lighthouse pilot (fig. 2) provides a universal media player allowing video, audio, PDF, presentations, and office documents to be viewed using a common user interface over multiple platforms and allowing tutor and student annotations on any kind of resource. This is provided for its pedagogic value and to maintain a consistent user experience, but in addition allows the collection of detailed educational trace data.

Whereas most learning platforms provide document-view or download levels of analysis, the Lighthouse pilot provides much more detailed learning analytics of what was viewed *within* a resource. In the player a summary is presented as part of the academic view of a teaching resource including a profile of which pages, or segments of video, have been viewed (fig. 3).

With these analytics it was immediately obvious, as I had always suspected, that a large proportion of student 'views' of a resource never progress beyond the first page of a document or few seconds of video. Whereas page-view statistic gave the impression of widespread student engagement, in fact many students were doing little more than click through and move on.

This level of knowledge was helpful both to give a sense of control and for directed pedagogic intervention. On one occasion I was able to see that most students had read (or at least looked at the pages of) a journal paper, but few had got more than 20-30% of the way into the paper. I was able to advise them that, although it was hard to read a full research paper, they could skip to the last section in which the theoretical techniques introduced earlier in the paper were used in practice by product designers [10].

While already useful, there is a lot more data available than is currently surfaced in the interface. As part of an investigation into future requirements, experimental drilldown graphs were produced of the data on my own modules using offline data extraction, and visualisation. I analysed these individually partly to see if this level of drill-down would be useful in its own right, and partly to work out if there were patterns of detailed behaviour that could be extracted to create different forms of summary view, or create pedagogically meaningful features. The latter would be useful to feed into systems like Purdue Signals [1] that perform data mining combining student behaviour and outcomes in order to create diagnostic aids.

Some reading patterns were expected (fig. 4) such as giving up part way through, or skimming and then reading parts in detail. However, other patterns emerged (fig. 5) such as skipping to the end and skimming backward (find the conclusions). Crucially, I was someone who both (i) understood the data analysis; and also (ii) understood the document and its teaching context. This allowed me to make sense of the places where students paused while skimming (usually the pictures), or make guesses at the reasons for giving up at various points [9].

I am currently working with other academics from different institutions and disciplines in order to determine what application elements or data features would be generally useful. However, this process is beginning with a far greater level of insight than would have been possible from uninformed data analysis. For example, Carmichael and Miller describe problems of interpretation even when higher education researchers study in further education contexts [4].




Figure 4: Some expected patterns (from [9], ©Talis)




Figure 5: Examples of actual student reading patterns. A lot more complex than fig. 4. (image ©Talis)

Forcing Focus

The final vignette is one where practical exigencies highlighted a profound gap in empirical and theoretical understanding. To some extent this is something that I've experienced throughout my research career in HCI; the need to deal with real user experience often exposes fundamental questions in computing or psychology. Here it is more that a product-focused context highlights gaps in the understanding of user experience.

Assuming that learning analytics can offer actionable insights, one question is how to notify these to students and academics. In studying this, it became clear that there is very little empirical data about the academic life. As an academic I was able to use my own anecdotal knowledge to make suggestions, but little hard evidence [8].

Indeed earlier in this paper I noted in passing that teacher motivation was important. This seems a relatively uncontentious statement, but even it is hard to justify empirically. Han and Yin [13] produced a recent comprehensive review of teacher motivation, but the vast majority of the 117 papers they reviewed relate to primary and secondary level teaching only.

Looking more widely at the pattern of actual (not normative) activities in an academic day, week or year there seemed to be little known. Possibly this is because academia is so varied it is hard to study, yet classic studies of knowledge workers have often dealt with small a diverse groups [15, 18]. It could just be that academics are worried about studying their colleagues for fear that they are seen as part of the 21st century management turn.

Conclusions

In these vignettes of academic and commercial researcher interactions with educational practice, we have seen both similarities and differences.

The most notable difference is that I found myself and reviewers adopting higher ethical standards for commercial research than for comparable academic work. Is this appropriate? There are conflicts of interest, but equally there are for the academic researcher, especially when experimenting on one's own students. For reviewers of commercial papers perhaps they should ask, "would I say this if the system were academic"; for academics in educational research perhaps they should ask "how would I behave if this were a commercial system I was testing.

On similarities, both teaching and the need for concrete commercial results push research boundaries. Fuzzy statements stand up neither to students' nor developers' scrutiny; one is forced to move from problems and questions to solutions and action. Furthermore we saw the power of being both researcher and educator: combining knowledge of both data processing and data meaning.

We also saw that students know a lot about technology, experience from which we can learn (even though they are far from infallible), and, perhaps most surprising, that despite extensive work on educational technology in universities, we know little about day-to-day academic life.

Acknowledgements

Many thanks to my colleagues at Talis, who put up with me as resident and usually whinging academic.

References

1. K. Arnold. 2010. Signals: Applying academic analytics. *Educause Quarterly*, 33, 1. <http://www.educause.edu/ero/article/signals-applying-academic-analytics>
2. BERA. 2011. *Ethical Guidelines For Educational Research*. British Educational Research Association. <https://www.bera.ac.uk/researchers-resources/publications/bera-ethical-guidelines-for-educational-research-2011>
3. S. Brewster. 1994. *Providing a structured method for integrating non-speech audio into human-computer interfaces*. PhD Thesis, University of York, UK. http://www.dcs.gla.ac.uk/~stephen/papers/theses/Brewster_thesis.pdf
4. J. Carmichael, K. Miller. 2006. The Challenges of Practitioner Research: Some insights into collaboration between Higher and Further Education in the LfLFE Project. In *Proc. of 3rd International CRL Conference (2005)*, Glasgow. 700–702. ISBN 1-903661-72-2. <http://www.lancaster.ac.uk/lfife/publications/pubsdocs/THE%20CHALLENGES%20OF%20PRACTITIONER%20RESEARCH.doc>
5. A. Dix, G. Ellis. 1998. Starting simple: adding value to static visualisation through simple interaction. In *Proc. AVI '98*, ACM, 124–134. DOI: 10.1145/948496.948514
6. A. Dix. 2002. Beyond Intention – pushing boundaries with incidental interaction. In *Building Bridges: Interdisciplinary Context-Sensitive Computing*, Glasgow University, 9 Sept 2002. <http://alandix.com/academic/papers/beyond-intention-2002/>
7. A. Dix, J. Finlay, G. Abowd, R. Beale. 2004. Modelling Rich Interaction. Chapter 18 in *Human-Computer Interaction*, third edition.
8. A. Dix and J. Leavesley. 2015. Learning Analytics for the Academic: An Action Perspective. In *Journal of Universal Computer Science (JUICS)*, 21(1):48–65. DOI: 10.3217/jucs-021-01-0048
9. A. Dix. 2016. Challenge and Potential of Fine Grain, Cross-Institutional Learning Data. In *L@S: Third Annual ACM Conference on Learning at Scale*, ACM, pp. 261–264. DOI: 10.1145/2876034.2893429
10. A. Dix. 2016. More than one way to flip a class: learning analytics for mixed models of learning. *Compass: Journal of Learning and Teaching*, 8, 12. DOI: 10.21100/compass.v8i12.275
11. A. Dix. 2017. Activity modelling for low-intention interaction. Chapter 7 in *The Handbook of Formal Methods in Human-Computer Interaction*. B. Weyers, J. Bowen, A. Dix, P. Palanque (Eds.), Springer.
12. W. Gaver. 1993. How do we hear in the world? Explorations of ecological acoustics. *Ecological Psychology*, 5, 4: 285–313.
13. J. Han, H. Yin. 2016. Teacher motivation: Definition, research development and implications for teachers. *Cogent Education*, 3, 1: 1217819. DOI: 10.1080/2331186X.2016.1217819
14. E. Helsper, R. Eynon. 2010. Digital natives: where is the evidence? *British educational research journal*. 36, 3: 503–520. DOI: 10.1080/01411920902989227
15. A. Kidd. 1994. The marks are on the knowledge worker. In *Proc. CHI '94*. ACM, 186–191. DOI: 10.1145/191666.191740
16. M. Prensky. 2001a. Digital Natives, Digital Immigrants : Part 1. *On the Horizon*, 9, 5: 1–6.
17. M. Prensky. 2001b. Digital Natives, Digital Immigrants Part 2: Do they really think differently? *On the Horizon*, 9, 6: 1–6.
18. S. Whittaker, D. Frohlich, and O. Daly-Jones. 1994. Informal workplace communication: what is it like and how might we support it?. In *Proc. CHI '94*. ACM, 131–137. DOI: 10.1145/191666.191726