



AUTHOR(S):

TITLE:

YEAR:

Publisher citation:

OpenAIR citation:

Publisher copyright statement:

This is the _____ version of an article originally published by _____
in _____
(ISSN _____; eISSN _____).

OpenAIR takedown statement:

Section 6 of the "Repository policy for OpenAIR @ RGU" (available from <http://www.rgu.ac.uk/staff-and-current-students/library/library-policies/repository-policies>) provides guidance on the criteria under which RGU will consider withdrawing material from OpenAIR. If you believe that this item is subject to any of these criteria, or for any other reason should not be held on OpenAIR, then please contact openair-help@rgu.ac.uk with the details of the item and the nature of your complaint.

This publication is distributed under a CC _____ license.

BROOKES eJOURNAL OF LEARNING AND TEACHING

PROMOTING GOOD PRACTICE IN LEARNING, TEACHING AND ASSESSMENT IN HIGHER EDUCATION

A practitioner reflection on teaching computer ethics with case studies and psychology

Vol. 7 - issue 1 - July 2015

Authors

Dr Michael Heron

Pauline Belford

Abstract

Three issues complicate teaching computer ethics in an undergraduate course. The first relates to the often technically intensive knowledge required to fully understand the complexity of real world examples. The second relates to the pedagogic expectations of students who see ethical and professional issues as of little importance to their eventual degrees. The third revolves around the fact that the official accreditation that is required of many professions is not mandatory for computing professionals, and so professional codes of conduct are optional. In this reflective discussion, we discuss these issues and the approach we have taken to resolve them. Our philosophy for teaching computer ethics revolves around the use of social psychology to illuminate the importance of the topic, and case-studies to simultaneously lower the burden of technical expertise while also incorporating hooks for the discussion of real world incidents. We discuss several psychological studies which inform our discussions, and the way in which they are delivered to overcome initial student objections to the material. We then discuss both the Case of the Killer Robot and the Scandal in Academia as case studies appropriate for inclusion in most undergraduate and postgraduate courses on ethics and professional issues.

Introduction

Awareness of relevant professional and ethical issues is an important part of many courses of higher education. For some professions, the accreditation of a formal body is a pre-requisite for being able to practise and as such the need to adhere to a code of ethics is something that can be easily understood by students. Within the field of computing however there are no professional accreditations required before one can practise and so these must be 'sold' on the basis of the perceived benefits available from membership. Accreditation, both individually and institutionally, from organisations such as the British Computer Society (BCS), the Institution of Electrical and Electronics Engineers (IEEE) or the Association of Computer Machinery (ACM) is perceived by many to be a valuable professional signifier (Williams, la Fosse, Harris & Thilthorpe, 2011; Denning & Frailey, 2011) and source of networking opportunities. Institutional or course accreditation allows for organisations to facilitate access to membership levels for graduates as well as codify a philosophy of continuous quality assurance (Durrani, 2011; Godweber, Davoli, Little, Riedesel, Walker, Cross & Von Kinsky, 2011; Impagliazzo & Gorgone 2002). Students and their parents in turn perceive this as a valuable benefit of choosing one university over another. This in turn creates a pressure on universities to offer courses which meet these accreditation standards.

To offer such accreditation a university must meet the requirements of the professional body. These requirements will almost always come with a need to inculcate relevant professional and ethical values into graduates. The BCS requires ethical instruction as a core part of accreditation (British Computer Society, 2012). The joint task force of the ACM and the IEEE likewise stress the importance of teaching ethical issues within their guidelines for undergraduate accreditation (Joint Task Force on Computing Curricula, 2013).

However, while the prospect of professional accreditation often appeals to students, the reality of the instruction often represents a sticking point (Cajander, Daniels, McFermott, von Kinsky, 2011). The technical requirements of the guidelines usually enjoy a tight match to student expectations, but this is rarely true of teaching professional values. A high level of institutional variation is to be expected, but within the field of computing and related topics it's often the case that ethical and professional issues are seen as less important than topics that teach practical skills (Makasiranondh, Maj & Veal 2011).

The authors of this paper have taught modules that are either fully or partially focused on professional issues at all levels of the higher education curriculum, from first year undergraduate to professional Masters degrees. For the latter, the module was delivered to serving police officers working with forensic computing and cybercrime and as such the importance of ethical context in their professional lives was already well understood.

For all other modules a certain degree of 'pushback' was expected and observed, where students (even for electives) questioned the relevance of the issues to their expected qualifications. A non-trivial portion of teaching time then is spent justifying the topic rather than exploring its complexities. Along with this issue goes the difficulty in relating theoretical instruction to real life. There are many examples of professional misconduct and ethical complexity to be drawn from the news and

academic literature, but in a field as fast-moving and complex as computing they can be difficult to isolate and generalise to illustrate wider trends.

In this paper, the authors reflect upon several years of intensive teaching of ethical issues at undergraduate and postgraduate levels, and discuss both the techniques used and the results obtained. The authors argue that underpinning ethical discussion with examples drawn from the psychological literature helps provide context for the importance of the topic, and that the use of formal case studies helps contextualise analysis of the issues.

Specific issues in teaching computer ethics

Within computing three key issues complicate the teaching of ethics and professional conduct.

The first of these is that computing is a fast-moving and highly technical field, and as a consequence there are certain pre-requisites of understanding for the implications of professional or ethical misconduct to be fully appreciated. A similar argument could be made for the complexity of medical ethics (Sugarman & Sulmasy, 2010) or legal ethics (Jacobowitz & Rogers, 2013) – it is not an issue unique to computing, but one that must still be addressed within the topic. For first year undergraduates in particular, professional issues must be discussed before any detailed technical expertise has been built. The common vocabulary of the field may not be shared, and the experience needed to understand the subtle implications of real-world incidents is rarely present.

Consider for example the widely discussed systems design failure of the Mars Climate Observer (McCurdy, 2013; Kumar, Raghu & Kumar, 2013). To understand the professional issues that led to the crash of the probe into the Mars atmosphere one must understand the foundational elements of systems analysis and design, the realities of cross-team development, and the implications of testing at the boundaries. These are not insurmountable issues, but they do complicate the job of finding relevant real life examples and presenting them to students. The problem becomes less pronounced the further up the modules are in the degree structure, but they never fully go away. Even students within the last stages of their honours degree understand, in depth, only a subset of the technical issues across the field of computing. The case of the Nanniebot (Sumner, 2011) as an example requires a reasonable grasp of the state of the art in AI 'chatterbot' technology, and unless that topic has been covered as part of the course the knowledge cannot simply be assumed. Both of these examples would be valid topics for discussion within a course on professional and ethical issues, but the pre-requisites for understanding are rather high. The job of finding suitable examples for instructional discussion is complicated and may not be generalizable outside of a single course or institution.

The context of these individual scenarios is important because best practice in pedagogy is to ensure participation rather than silent absorption of material, and meaningful participation in a discussion about an ethical issue requires informed viewpoints that can be expressed, challenged and debated. That in turn requires students who can insightfully assess a scenario from multiple angles. In the best discussions about ethical and professional issues, personal experience informs the

analysis, which means that the wider the pool of opinions that can be solicited, the better the following discussion. When ethical scenarios require the instructor to provide the context, students are largely limited to discussing the issue from the viewpoint that was provided. Morality and ethics are often highly individualised (Heron and Belford, 2014) and meaningful analysis only occurs when competing viewpoints are compared and contrasted.

The second issue relates to the pedagogic expectations of computing students. The study of ethics and professional issues is widely considered to be dry and irrelevant. They have little time, in the main, for philosophical definitions or arguments. They have little interest, on the whole, in nuanced hypotheticals. They believe, more often than not, that ethics is a simple issue of common sense and that they are generally 'good people'. For many 'day to day' ethical 'dilemmas' this might even be true. Most students, after their first few weeks of university instruction, could give a perfect description of what constitutes plagiarism. However, our direct experience in investigating plagiarism over the years has shown that their theoretical appreciation is not necessarily matched by their practical understanding. Especially in disciplines such as programming, the line between what is and is not acceptable may be hard for novices to discern. In many respects, all programming is plagiarism – there are only so many ways to write a loop or construct a simple class, and all that follows is variation on a theme. Programming too inculcates, at its core, an appreciation of the value of **reuse of code**. It is not in the obvious cases where ethical decision making is complex, but instead in the real-world **edge cases**.

Thus, the educator must fight against the perception, widely held, that ethics are 'easy' and that we can all be trusted to do the right thing. The educator too must work to ensure that student understanding is on a firm theoretical basis so that the context of an activity can be properly evaluated as to its relevance to other situations. Generalising with respect to context is a powerfully important skill. A series of discussions about real-life incidences is not enough without providing the necessary underpinnings from which students can extrapolate.

Finally, the professional accreditation that is an important part of many careers is entirely opt-in within computing. One does not need accreditation to create a website, administer a database or write a computer program. Students and staff sometimes see progressing through a module designed around the requirements of these optional professional bodies as futile busywork.

Where the modules are electives, they are often taken as an 'easy option' by students who wish to shape their courses into being less technically intensive. It is rare, in our experience, that a class consists of those who have a deep desire to learn the topic. We have found too that guidelines for adherence to a professional code that students may have no deep commitment towards makes for dry content. Membership of professional bodies in computing is reasonably high but not absolute and many students will not take the trouble to seek out membership at even the lowest grades. Accreditation requires respect for an organisation's ethical codes, but students often see this as completely tangential – if they wish to join the organisation, they will make an effort then to learn the codes to which they must adhere.

None of these issues are unique to computing, but they do come together in such a way as to frustrate attempts to teach the topic. Computing is in many ways an applied subject area where

students learn how to master technical content. This separates it in tone and expectation from, for example, law degrees where rhetoric is prized and medical degrees where legal implications are significant and understood to be important.

An approach to teaching ethical and professional issues in computer science

Within the ethical courses we have taught, we have followed a certain blueprint that allows for us to make overcoming the issues outlined above a core part of the module. To deal with the issue of contextualising real world examples we make use of a formal case study. This provides a scenario that is suitably complex whilst also accessible enough that those without deep technical knowledge can still meaningfully discuss the implications of the actions taken by the various protagonists. In previous years we have heavily used the Case of the Killer Robot (Epstein, 1994a; Epstein, 1994b; Epstein, 1996), but its age is such that it is no longer entirely relevant to modern computing. Instead, a bespoke scenario called The Scandal in Academia (Heron and Belford, 2014) was written as a 'spiritual successor'.

While there are technical aspects to the case study, they are delivered through the lens of a newspaper reporting on a growing academic scandal and thus our difficulty in truly understanding the workings of complex research algorithms plays a key part in the delivery of the scenario material (Heron and Belford, 2014). Within the scenario, we are given few technical details but this becomes an entrance point to a discussion about the replicability of research and the difficulty of truly understanding a technical ethical issue when we have no deep knowledge of the topic. Both case studies will be discussed more fully later in this paper.

The issue of pedagogic expectation is handled through directly challenging the core pre-suppositions of those in the classroom by incorporating findings of social psychology. The first question asked of students is 'How many of you feel that, on balance, you are good people?' Within most classes, a healthy majority of around 70-75% of the hands will go up, with the rest being either uncommitted or jokingly resistant to the idea.

The next question becomes 'How many of you think you could kill someone for a good enough reason?' Fewer hands go up, usually around 30-40%, and this becomes an introduction to a short discussion on what constitutes a 'good enough reason'. Views expressed here range quite widely, from stridently reactionary to equally strident progressiveness.

The third question becomes 'How many of you think you would kill someone just because a person in a white coat tells you to?' Usually, no hands go up at this point, which leads to the statement 'Statistically speaking, two thirds of you are wrong'. This then leads into a discussion of the Milgram Experiment (Milgram, 1965; Blass, 1999), a well-replicated study showing the power of obedience, authority and context. The Milgram Experiment leads into a discussion of the Stanford Prison Experiment (Zimbardo, 2007) and then into the Asch Line Conformity experiments (Asch, 1956). All

three of these studies show the incredible power that social contexts can have on what we consider to be normal ethical behaviour and how someone doesn't have to be evil to perform evil actions. These will be discussed more fully in the next section.

This technique, while effective, is not to indicate that the psychological literature is key to understanding computer ethics. Instead it is to reinforce the idea that ethical study is relevant to their professional development by nullifying the argument that it's just a matter of common sense. The intention is to reinforce the idea that there are rarely such things as 'good people' and 'bad people', but instead 'people in different social contexts'. This rapidly disabuses classes of the idea that they as individuals can simply be trusted to do the right thing. It also serves a useful purpose for setting part of the theoretical context of an understanding of ethics – by examining the psychological basis of what might be considered 'evil actions', we rapidly discover that ethics, morality and professional conduct are more about social psychology than they are about philosophy. We need not follow through the philosophical reasoning of the great philosophers. While there is much of value in this area, it is difficult to inspire students with abstract logical reasoning if they are not already philosophically inclined.

Since these studies draw heavily from the psychological literature rather than the computing literature, the experiments are usually new to the students. The implications of the experiments are likewise somewhat shocking, and this serves as the jolt necessary to overcome initial resistance based on an expectation of what instruction on professional issues will involve. It also highlights that ethics need not be about academic abstractions or hypotheticals. The studies may not be full of hard data or quantifiable measures, but they reveal more than simple discussion or conjecture. They emphasise that understanding the importance of ethical and moral behaviour is more than simply talking about what other people have done, but considering what we might ourselves do if the social environment in which we function was different.

The last issue, that of the fact that professional codes may not be relevant to those who do not intend to join a professional body, is addressed by using an ethical code as a lens by which to evaluate the actions of others. The ethical code does not become a topic of passive instruction, but instead a tool for active evaluation. This in turn makes it possible for teaching material to be flexible and applicable to what may be subtly incompatible codes. 'With reference to the BCS code of conduct, discuss the ethical and professional implications of X's activities in the following scenario'. Thus, while the code must be discussed within the module its true weight is only tangible within the assessment activities. With all due respect to the work that went into creating these codes, it is difficult to interest students in a line-by-line analysis of the various guidelines. Such codes are best used as a tool, rather than a topic. This in turn has the benefit of making each of the guidelines something that gets scrutinised for its relevance – students give due consideration to nuance and the wider implications of the guidelines because they are the evaluative instrument which helps them gain the module grades they want.

Psychology and ethical issues

The use of a case study as a tool to bind together disparate elements of ethical and professional issues in a computer science course is uncontroversial. Pedagogic good practice argues that case

studies are often relevant (McDade, 1995) and the ones that have been discussed in this paper have been used with considerable success. Grounding the discussion in the social psychological literature however may be more complicated. Generally, computing professionals will not be fully conversant in the psychological literature, and mastering a new discipline full of nuance and complexity may not be possible within institutional constraints and reasonable workloads.

However, psychology and computer science have always had numerous areas in which they overlap – particularly in areas such as usability and artificial intelligence. There are various well-replicated findings in the social psychological literature that illuminate modern issues of computer ethics as well as help overcome the initial resistance that students may feel towards the topic. There are numerous online videos that show the experiments and their implications and these can be used to confront students with the reality of psychology. These have the added benefit of often being more effective than a simple description of the experiment and the findings. Seeing how participants within the Milgram experiments squirm and prevaricate before eventually progressing through the experiment illustrates the internal tensions better than any lecture content could. For many of these experiments, the findings are so shocking that only direct exposure to the studies really convinces.

Consider the key element of the Milgram Experiment. The experiment was repeated in many variations, but in the best known version 65% of research participants delivered what was likely to be fatal electric shock to someone they thought was another research participant simply because they had been repeatedly instructed to do so by a man in a white coat. The Milgram Experiment sets the context of the power of authority even when such authority is transient or illusory. However, relating it directly to students illustrates only one part of the important message – that many of us, under the right circumstances, could be compelled to do terrible things. It communicates that the 'right circumstances' need not be extreme or even outside our every-day experiences, but it shows only the impact of social context on those who are its victims. It shows how authority can be misused, but the findings relate to the one subject to authority rather than the one wielding it. Milgram's own documentary *Obedience* (1965) includes much useful and harrowing footage.

The Stanford Prison Experiment on the other hand is useful in exploring the other side of the equation. Within the Stanford Prison Experiment, twenty-four male students were selected as participants from a volunteer pool of seventy-five. All participants were screened for psychological, medical and physical issues and then randomly assigned to one of two experimental conditions – half were selected as guards and half were selected as prisoners. The prisoners were then picked up by police car from their homes, handcuffed, strip-searched and deloused, and then placed in a replica prison built within Stanford University itself. The guards were given a set uniform consisting of khaki shirts, truncheons and shades. They were then told their job was to control the prisoners.

The experiment was scheduled to last two weeks, but ended after only six days because of a constant escalation of authoritarian measures enacted by the guards to control the prisoners, including psychological and physical harassment, and the turning of prisoners against each other through victimisation and the encouragement of peer pressure. While the study has numerous critics, and its findings cannot be easily generalised, it shows at least one perspective as to how quickly 'normal' people can become brutally sadistic. The random assignment of individuals to guard and prisoner conditions offers some protection against a self-selection bias. It suggests strongly that

we are all susceptible in a scenario where we are given power, little oversight, and an adversarial environment where we are only one of a number charged with retaining control. Seeing the footage from the experiment, along with Zimbardo's commentary, allows for the study to be incorporated into a course without necessitating expertise in social psychology. The documentary (Zimbardo & Musen, 1989) contains much useful content.

Within the Asch line conformity experiments, we see the difficulty of standing out in a crowd. As with the Milgram Experiment, the participant in the research study is a patsy, and is placed in a room with seven 'confederate' allies of the researcher. They are shown a series of three lines, and their job is to say which is the longest, with the patsy being the last to answer in all conditions. For the first few charts, all confederates are told to reply correctly. As the experiment continues, confederates are asked to declare the same wrong answer to see what happens to the answers provided by the patsy. In the control condition, where no collusion amongst confederates occurs, the error rate of the patsy is less than 1%. In the experimental group, 75% of participants give a wrong answer to **at least** one of the questions, known in the literature as a 'conformity rate'. In total, a third of the answers given were wrong. The lesson here is that it is difficult to speak out when there is a wall of disagreement, and that instead many people will simply 'go with the crowd' to avoid standing out. However, many variations of this experiment have been conducted, and shown that if even **one** confederate goes against the rest of the group then the patsy is emboldened and conformity rates plummet. The experiment shows that peer pressure is not simply a schoolyard phenomenon and there are difficulties in being the 'outsider voice' in even a simple scenario where we are absolutely sure we are right. However, it also shows that if we speak out we are likely not alone in our views—it's hard to be the first, but it is likely that being the first will mean that we will not remain alone.

No psychological study of the importance of Milgram, Stanford or Asch escapes their long decades of scrutiny without being questioned, however – students should be encouraged to query the findings, the methodology, and the interpretation of the results. Perry (2013) for example argues that the results of the Milgram experiment should be dismissed because the 'researcher' within the study was actively prompting participants for compliance. There are reasonable questions to be raised here as to whether that makes a difference to the key findings. Perry also argues that most participants 'saw through' the whole thing based on post-study debriefing – here a class can meaningfully discuss the nature of recall, false recall, social desirability bias and the hindsight effect. For the Stanford Prison Experiment, other studies (such as Haslam and Reicher, 2006) fail to replicate the results, but examination within class of the methodological differences in the two studies can offer considerable insight into the conditions under which societal pressure can be brought to bear, or not. The key benefit to be had in deploying these experiments is to build student buy-in. We need not insist they are unquestionable or even that they are correct – we simply need to ensure that students have suitably discussed the issues and implications that lead to that judgement.

Other valuable tools exist to deconstruct the idea that we are wholly responsible for our actions. The bystander effect demonstrates that diffusion of responsibility can create situations in which passive inaction leads to tragedy (Fischer et al, 2011). The reciprocity effect (Falk & Fischbacher, 2006) explains much of the viral property of social network engagement and beyond, eloquently summed up in Sheldon's complaint in the television sitcom *The Big Bang Theory* that, "You didn't get me a gift! You gave me an obligation!" The concept of deindividuation (Diener, 1980) addresses why

discussion on the internet can be so toxic on controversial topics such as discussed in Heron, Belford and Goker (2014). Social proofing (Sherif, 1935), or the human strategy of looking to those around us to see how we should behave in uncertain situations, underlines memetic propagation and the driving power of Internet bandwagons. All of these studies have implications that directly inform the teaching of computer ethics and provide an academic weight to the topic that is grounded in empirical evaluation.

Case studies

By the far the most important tool we have when teaching ethics is a coherent case study. A good case study offers a way to reduce the technical knowledge required of students while also showing the wider implications of professional and ethical misconduct. A good case study is layered, introducing new complexity and revisiting previous judgements to see if they still hold true in the light of new data. We live in a world of imperfect information, and it is rarely the case that we will have all the facts available to form a flawless impression in any given scenario. Instead, we must work with limited facts, weight them accordingly, and then arrive at an opinion based on what we know. With scenarios drawn from real life, we rarely get any further insights into the nuance of the incident once the initial interest has worn off – we are rarely given an opportunity to revisit our opinions and have them challenged by new revelations. A well-designed case study can permit that to happen.

Within our teaching we have used the Case of the Killer Robot very heavily and with considerable success. It has been a staple of our modules on ethics and professional issues for over a decade, and as such we have experimented widely over the best way to handle it. The case study consists of a core of nine newspaper articles for a fictional incident in which a robot kills its operator as a result of programmer incompetence. The programmer is indicted for manslaughter, and the newspaper then charts the gradually evolving story. We find out about the culture of fear at the company, the slipping standards for 'correctness' of software, the institutional politics, personality conflicts, faked test data, insufficient training, and so on. The study gradually layers these revelations on top of others until we end up with a large cast of characters, all of whom have a role to play in the death of the robot operator. By the end, the best we can say is 'Nobody is to blame because everybody is to blame', but we begin with a very clear villain – the programmer who wrote the code that malfunctioned.

Each of the nine revelations is delivered within our classes in isolation. We treat each class as if the new article was just published, and then the lecture content of the module is an overview of how it all relates together and where real life has inspired the content. The Case of the Killer Robot drew heavily from real life incidences, but fictionalised them and used the modified versions as a kind of 'unit of ethical instruction' that could be incorporated into the on-going narrative. As such, each article brings with it a range of very interesting discussions about, for example, programmer liability in the development of computer software or real world incidences of deadly malfunctioning robots. The case study itself was inspired by the infamous THERAC-25 incident (Epstein, 2010; Zeikowitz, 2012). Each class within the module is focused on only one of the articles, which builds two key things into the class discussions:

- Natural pauses in the exploration of the case study, which means students have time outside the class to contemplate the story so far and to solidify their judgements before they are challenged by later story elements.
- A requirement to explore the full depth and nuance of the newspaper articles, avoiding shallow judgements.

Previous modules using the case study, primarily those where it is only a portion of the module rather than the key educational vehicle, have experimented with providing 'packages' of articles. However, in these accelerated variations the discussions tend to be less edifying because of the need to progress onto the next – contemplation is truncated, judgements do not get the necessary time to settle, and new revelations have diminished impact because we have not yet fully internalised the previous revelations.

Along with the discussion of the case study itself, students are encouraged to build a map of relationships between the cast of characters and the incidents that occur. By the end we find this map is huge and full of subtleties that the text does not explicitly outline. We see at a glance which characters are peripheral and which are deeply embedded in what happened, and along the way we need to consider the way in which the actions of one participant impacted on the outcome of the whole incident. By requesting students 'start afresh' each time, we get an evolving diagrammatic overview of the growing complexity of the case study – early charts have four or five major participants. Later ones may have thirty or more, with a complex web of relationships. By comparing the current chart against the previous chart, we are presented with a high level view of what our new article has told us about the context of what has happened.

Short lectures linking to real world scenarios, inter and intra group discussions and the development of organisational charts are the default activities of these classes, although several variations have been adopted.

However valuable the Case of the Killer Robot has been though, it is now at the point where its age begins to show. Contrary to helping deal with the knowledge deficit of technical issues, several parts of the case study require a short lesson in historical context – controversies that once raged wildly during the 90s have long since been resolved. Technologies and techniques mentioned in the study have faded into obscurity and are only occasionally taught. Articles purporting to relate to an analysis of modern computer interfaces are anachronistic. It remains a valuable and fascinating case study, but it is hard to convince someone who has not seen its use first-hand that it is worth adopting because of its age.

Having decided that too much time was being spent explaining the historical context of the case study, a spiritual successor to the Case of the Killer Robot was written. It's called 'A Scandal in Academia' (Heron & Belford, 2014) and adopts the same style of delivery but with a brand new scenario that incorporates many ethical issues that modern technology has introduced. These include the issue of permanence of online information and the dangers of social networking. It also emphasizes numerous issues that were less prominent during the nineties – the issues of accessibility, misogyny and corporate influences are important parts of the Scandal in Academia.

The story in the Scandal of Academia is that a prominent university academic, Professor John Blackbriar, is suspended by his institution as a result of allegations of academic misconduct. As the case study progresses, we find out that he has been 'smoothing' his research data by subjectively identifying 'outliers' that worked against his previous conclusions, and then publishing his results based on what the newspaper reporting on the allegations terms 'fake data'. Each newspaper article introduces new ethical issues, within the context of an evolving crisis at a fictional university, and links these back to previous revelations. The Scandal in Academia consists of fourteen parts, with the intention that these are delivered in the same way one might use the Case of the Killer Robot. An illustrative discussion of the first two articles in the case study is available (Heron & Belford, 2014) to show the academic context that may underpin use of the study within undergraduate courses.

The Scandal in Academia has been trialled with second and third year students taking a module on ethics as part of their undergraduate computing courses. The intention of these trials was to assess its applicability as a replacement for the Case of the Killer Robot. While it is inappropriate to generalise based on the subjective view of the author, the discussions that followed the classes were marked by how much more inclusive they were in terms of viewpoints expressed. The Case of the Killer Robot suffers somewhat in focusing on a reasonably technical context and many of the students with whom we have used it do not self-identify as programmers, software designers, or any of the other user archetypes that emerge throughout the story. It becomes difficult for students to truly empathise with the protagonists of the Case of the Killer Robot because both their current context and presumed career trajectories are widely divergent.

However, within the context of a university course everyone has experience of academics. Students are aware of the importance of research to many universities, and can empathise with the plight of postgraduate students and undergraduates buffeted by the political and reputational winds of the Higher Education institute they attend. The Scandal in Academia, judged on its initial informal exposure, would seem to offer more opportunities for non-technical students – the ones who tend to opt-in to elective modules on ethical and professional issues – to appreciate the relevance of the struggles the characters undergo.

The Scandal in Academia was used in two ways – the second year students were given the study in a phased delivery model, where articles were delivered one at a time and then discussed in groups. No formal lecture discussion of the case study was provided for these – the intention was to explore the views expressed by students without prompting or contextualising. A wide range of conflicting views were expressed including many that the authors of the case study had not considered.

The third year students were also given the article content in a phased delivery model and asked to discuss it in groups. For the final class on the case study they were separated into 'prosecution' and 'defendant' and told that their job was to argue for the culpability, or otherwise, of Professor John Blackbriar. They would be given ten minutes to make an opening statement, then ten minutes to refute the arguments of their opponents, and then ten minutes to make a closing statement.

Students were allocated to groups, and not allowed to select their own. This adversarial model required many students to reconsider the evidence from an alternate perspective and this in turn required them to consider counter-arguments against both their own initial beliefs and those that they

expected their 'competitor' group would express. Thus, rather than simply focusing on their own conclusions they were forced to consider those of an alternate perspective and find the best way to nullify them in debate.

Overall, student responses to the updated case study were positive, with several students asking if they would be able to get a copy of the study to consider outside of class time. The reduction of time required to contextualise the study translates directly into more time discussing the ethical and professional issues at the core of the scenario. Its increased accessibility offers an opportunity for a wider range of perspectives to be addressed because students can see the relevance of the discussions to dilemmas they may face in their own lives.

However, this was not intended to be a full evaluation of an experimental trial – the impression gained as to the value of the case study is entirely subjective, and since one of the authors of the study also conducted the classes it is hard to argue convincingly that the impressions gained are unbiased. Whether one prefers the Case of the Killer Robot or the Scandal in Academia, it is the use of the case study at all that is a powerful tool for contextualising ethical and professional issues. The authors believe that, when used in conjunction with effective lectures aimed at relating these fictionalised examples to real life scenarios and the psychological literature, case studies offer the best mechanism for challenging students to fully engage with the topic.

Conclusion

Teaching ethics to students who are not inherently interested in the topic is challenging. First, we must convince them that ethics is a topic worthy of discussion, which requires us to overcome the impression that it is merely a matter of common sense. We must disabuse students of the belief that being a 'good person' is enough to protect against ethical and professional misconduct. We must examine how our actions are influenced by the actions of others. We must discuss the pressures of the social contexts in which we function. The social psychological literature, rather than philosophical or scattered real life examples, helps communicate this lesson.

Once the value of the topic is appreciated, there is more value in moving to a formal case study that offers an opportunity to explore in depth the ethical and professional issues encountered in day-to-day life. Both the Case of the Killer Robot and the newer Scandal in Academia are appropriate here, although the former requires a considerable investment of time in explaining the historical context in which the incident is located. Case studies offer an opportunity to explore issues divorced of irrelevant real-life context and complexities. They also permit for the activities of the participants to be revisited in light of future revelations in a way that is rarely possible for real world examples. The best case studies are heavily inspired by real life, and so the revelations contained within both the Case of the Killer Robot and the Scandal in Academia permit for the original examples to be brought in to reinforce the relevance and 'believability' of the narrative.

A case study then provides a context in which we can examine, in depth, the actions and motivations of agents within an ethically complex scenario. Coupling that case study to the psychological

literature permits for the discussions that follow to have a tangible weight without the abstraction or unwieldy terminology of ethical philosophy.

Linking professional ethical codes to this model of delivery is best done as an evaluative tool rather than as a topic of discussion in and of itself. Since membership of professional bodies is optional, and computing practitioners may be members of one, all or none, it is rarely appropriate except for the purposes of external accreditation to incorporate one specific code within the teaching content. The requirements of accreditation can be addressed more appropriately within assessment.

This is a reflective practitioner account of an approach to teaching computer ethics. As such the success we have personally had cannot necessarily be generalised beyond the environments in which it was used. However, the high quality of discussions along with the incorporation and facilitation of varied viewpoints argued persuasively for an adoption of this or a similar model when the formal delivery of ethical and professional content is required with a technical computing course.

References

Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9), 1-70.

Blass, T. (Ed.). (1999). *Obedience to authority: Current perspectives on the Milgram paradigm*. Psychology Press.

British Computer Society. 2012. *Guidelines on course accreditation: Information for universities and colleges*. [Available online from http://www.bcs.org/upload/pdf/hea-guidelinesfull-2012_1.pdf]

Cajander, Å., Daniels, M., McDermott, R., & von Kinsky, B. R. (2011, January). Assessing professional skills in engineering education. In *Proceedings of the Thirteenth Australasian Computing Education Conference-Volume 114* (pp. 145-154). Australian Computer Society, Inc..

Denning, P. J., & Frailey, D. J. (2011). Who are we—now?. *Communications of the ACM*, 54(6), 25-27.

Diener, E. (1980). Deindividuation: The absence of self-awareness and self-regulation in group members. *The psychology of group influence*, 209-242.

Durrani, T. S. (2011, January). Institutional accreditation United Kingdom model—Scottish variation. In *Institutional and Programme Accreditation (IWIPA), 2011 International Workshop* on (pp. 1-5). IEEE.

Epstein, R. G. (1994a). The case of the killer robot (part 1). *ACM SIGCAS Computers and Society*, 24(3), 20-28.

Epstein, R. G. (1994b). The case of the killer robot (part 2). *ACM SIGCAS Computers and Society*, 24(4), 12-32.

Epstein, R. G. (1997). *The case of the killer robot: stories about the professional, ethical, and societal dimensions of computing*. New York/Chichester: Wiley.

Epstein, R. G. (2010, March). Workplace Issues in an Undergraduate Software Engineering Course. In *Software Engineering Education and Training (CSEE&T)*, 2010 23rd IEEE Conference on (pp. 1-8). IEEE.

Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, 54(2), 293-315.

Fischer, P., Krueger, J. I., Greitemeyer, T., Vogrincic, C., Kastenmüller, A., Frey, D., Heene, M., Wicher, M., & Kainbacher, M. (2011). The bystander-effect: a meta-analytic review on bystander intervention in dangerous and non-dangerous emergencies. *Psychological bulletin*, 137(4), 517.

Goldweber, M., Davoli, R., Little, J. C., Riedesel, C., Walker, H., Cross, G., & Von Kinsky, B. R. (2011). Enhancing the social issues components in our computing curriculum: computing for the social good. *ACM Inroads*, 2(1), 64-82.

Haslam, S. A. (2006). Rethinking the psychology of tyranny: The BBC prison study. *British Journal of Social Psychology*, 45(1), 1-40.

Heron, M. J., & Belford, P. (2014). Ethics in context: a scandal in academia. *ACM SIGCAS Computers and Society*, 44(2), 20-51.

Heron, M. J., Belford, P., & Goker, A. (2014). Sexism in the circuitry: female participation in male-dominated popular computer culture. *ACM SIGCAS Computers and Society*, 44(4), 18-29

Impagliazzo, J., & Gorgone, J. (2002). Professional accreditation of information systems programs. *Communications of the Association for Information Systems*, 9(1), 3.

Jacobowitz, J. L., & Rogers, S. L. (2013). Mindful Ethics—A Pedagogical and Practical Approach to Teaching Legal Ethics, Developing Professional Identity, and Encouraging Civility. *St. Mary's Journal on Legal Malpractice and Ethics*, 4.

Kumar, C. S., Raghu, D., & Kumar, P. R. (2013). A Domestic Case Studies Probability to Overcome Software Failures. *Journal of Telematics and Informatics*, 1(1), 20-25.

Joint Task Force on Computing Curricula. 2013. *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science*. DOI: 10.1145/2534860

- Makasiranondh, W., Maj, S. P., & Veal, D. (2011). Student opinions on their development of non-technical skills in IT education. *Modern Applied Science*, 5(2), p3.
- McDade, S. A. (1995). Case study pedagogy to advance critical thinking. *Teaching of psychology*, 22(1), 9-10.
- Milgram, S. (1965). Some conditions of obedience and disobedience to authority. *Human relations*, 18(1), 57-76.
- Milgram, S. (1965). Obedience. *Milgram Films*.
- McCurdy, H. E. (2013). Learning from history: Low-cost project innovation in the US National Aeronautics and Space Administration. *International Journal of Project Management*, 31(5), 705-711.
- Perry, G. (2013). Deception and Illusion in Milgram's Accounts of the Obedience Experiments. *Theoretical & Applied Ethics*, 2(2), 79-92.
- Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology* (Columbia University).
- Sugarman, J., & Sulmasy, D. P. (Eds.). (2010). *Methods in medical ethics*. Georgetown University Press.
- Sumner, J. (2011). Teaching Source Criticism and Independent Investigation in the History of Science, Technology and Medicine. *Learning and Teaching in Philosophical and Religious Studies*, 10(2), 195.
- Williams, C., La Fosse, S., Harris, R., & Thilthorpe, A. (2011). The Relevance of Professional Bodies. *ITNOW*, 53(6), 12-13.
- Zelkowitz, M. V. (2012). what have we Learned about Software engineering? *Communications of the ACM*, 55(2), 38-39.
- Zimbardo, P., & Musen, K. (1989). *Quiet rage*.
- Zimbardo, P. (2007). *The Lucifer effect: Understanding how good people turn evil*. Random House LLC.