Measuring the Longitudinal Impact of Teaching Programming Ethics to First Year Students

Damian Gordon[1], Michael Collins[1], Paul Gibson[2], Dympna O'Sullivan[1]
[1]Technological University of Dublin (Ireland)
[2]Telecom Sud-Paris (France)
Damian.X.Gordon@TUDublin.ie, Michael.Collins@TUDublin.ie
Paul.Gibson@Telecom-SudParis.eu, Dympna.OSullivan@TUDublin.ie

ABSTRACT

In the second semester of the academic year 2020-2021 (January to May 2021), a group of 175 first year computer science students in the School of Computer Science at Technological University Dublin, Ireland, were presented with four case studies on programming-related ethical scenarios, which was assessed using a combination of continuous assessment and examinations as part of the Erasmus+ Ethics4EU project (O'Sullivan and Gordon, 2020). The goal of these cases was not to tell the students the "correct answer" to an ethical dilemma but rather to get them to explore in groups and reflect personally on what they would do when presented with a particular scenario, analysing what they think is the right thing to do. The approach to teaching this content represented a shift of teaching paradigm from Behaviourist approach to a more Constructivist approach (Ling and Ling, 2016), which was supported by a three-part lesson structure - (i) teaching by lecturers, (ii) discussion in break-out room by groups of students, and (iii) sharing key ideas discussed using a chat tool and the Padlet notice board. Based on feedback from the students, they found the lessons extremely enjoyable and engaging and they felt the break-out rooms were essential in allowing them to see the content through others perspectives. In this follow-up study six months later, the students were asked to recall the content they could remember from the previous lessons as well as the key ideas they learned from those case studies. The outcomes were as expected, i.e., some students have almost no recollection of the case studies but remembered general ethical themes, whereas others had detailed recollection of both the content of the classes and the key ethical issues discussed. This supports Grosz's notion of a "distributed pedagogy", where ethics needs to be infused throughout the Computer Science curricula to remind students of some key ethical ideas, as well as to give the students a better understanding of the ethical impacts and possible harmful effects of the technologies they implement (Grosz *et al*. 2019).

1. INTRODUCTION

The importance of teaching ethics as part of a computer science degree has been recognised for some time, but it wasn't until 1985 that computer ethics began to emerge as a separate discipline. This was the year that two seminal publications were produced, Deborah Johnson's book "Computer Ethics" (Johnson, 1985) and James Moor's paper, "What Is Computer Ethics?" (Moor, 1985). Johnson's book was the first to concentrate on the ethical obligations of computer professionals, and thoughtfully identifies those ethical issues that are unique to computers, as opposed to business ethics or legal ethics. In Moor's 1985 paper, he defined computer ethics as "*the analysis of the nature and social impact of computer technology and the corresponding formulation and justification of policies for the ethical use of such technology*", and he argues that computer technology makes it possible for people to do a vast number of things that it wasn't possible to do before and since no one could do them before, the question may never have arisen as to whether one ought to do them.

The field of computing ethics continued to evolve in the 1990s with the concept of "value-sensitive computer design". This was based on the notion that potential computing ethics problems can be avoided while new technology is under development by anticipating possible harm to human values and designing new technology from the very beginning in ways that prevent such harm (Flanagan, *et al*., 2008; Brey, 2012). At the same time, others including Donald Gotterbarn (1991), theorised that computing ethics should be seen as a professional code of conduct devoted to the development and advancement of standards of good practice for computing professionals. This resulted in the

development of a number of codes of ethics and codes of conduct for computing professionals. One important example is the ACM code which was first established in 1966 under the title "Guidelines for Professional Conduct" with the aim of upholding ethical conduct in the computing profession (Gotterbarn *et al*., 2018). That code of ethics has undergone various updates while keeping ethics and social impact as its main purpose. One of its most important updates was in 1992 when it was renamed to "ACM Code of Ethics and Professional Conduct" and was made up of 25 ethical principles for professionals to follow.

2. THE LESSONS

As mentioned previously a group of 175 first year computer science students were presented with four case studies on programming-related ethical scenarios, which was assessed using a combination of continuous assessment and examinations. The case studies were supported by a three-part lesson structure (i) teaching by lecturers, (ii) discussion in break-out room by groups of students, and (iii) sharing key ideas discussed using a chat tool and a Padlet noticeboard. The topics covered by the case studies are reviewed in the following subsections.

2.1. Irish State Examinations 2020

In 2020, the second level state examinations of Ireland had to be replaced due to the outbreak of COVID-19, and the system that was put in place was an estimated-grading system. The grades were estimated based on a number of factors, including the students' previous results, their school's performance, and an estimation from their teachers (Kelly, 2021). The algorithm to implement this process was implemented by an international IT company (based on a tendering process). As part of the quality assurance process, two separate professional bodies oversaw the development of this software, and checked the outcomes of the process. After the estimation process was undertaken and students were assigned their grades, these external bodies detected four errors which resulted in approximately 10,000 students being assigned a lower estimated grade in one or more subjects to what they should have received. As soon as the errors were detected, the affected students were identified, and corrections were made. However, the delays in making these corrections meant that some students had not received correct offers for university places and had to wait to commence their third-level study in the following academic year.

2.2. Search Engine Bias

There are a number of potential ethical issues with search engines (Canca, 2022), and in particular with the Google search engines, for example, when the user enters a search string, there is an auto-complete feature that offers the users suggestions on how to complete their strings. The algorithm for auto-completion uses a combination of the users' search history as well as common searches from other users, this can sometimes result in humorous suggestions, and sometimes result in offensive ones. More concerning is the fact that when a search string is entered, the ranking of the results is not simply based on which websites that the engine thinks most suits the search, but instead there are a range of additional factors that impact the ranking of results, include:

- *Fee Payment*: Individuals and organisations can pay a fee to Google to improve their ranking
- *Search Engine Optimisation*: Individuals and organisations can trick search engines using techniques such as adding additional metatag information to a web page that is hidden from the user but is used by the search engine when searching and ranking the page.
- *Political and Legal factors*: In certain jurisdictions, political and legal factors will influence the search results and subsequent ranking of results.

2.3. Judicial Sentencing Software

In the United States of America, some judges in court use commercial software to advise on sentencing based on predicted recidivism rates (Hillman, 2019). These systems have been investigated by

journalists working with Data Scientists, and have been found to have both racial and gender bias in their predictions. The systems tended to overestimate the potential for Persons of Colour to reoffend and underestimate Caucasians' potential for re-offense. Similarly, the systems tended to overestimate male defendants' potential to reoffend whilst making underestimations for female defendants. These systems should be unbiased in all scenarios but there is clear evidence that biased factors have been included by the system's programmers.

2.4. Automotive Accident Algorithms

There has been a significant growth in the development and commercialisation of self-driving cars, or those cars that require minimal human intervention to navigate between two geographical points (Lin, 2016). These cars use an array of sensors to capture data and make decisions using artificial intelligence (A.I.) to route-find and make navigation decisions. There is potential that an accident could occur that would result in the car having to decide who to save and who to let die, therefore programmers of these cars need to be cognisant of parameters that might include legal, moral, cultural, ethical and geographical factors.

3. THE FIRST SURVEY

Following the lessons, a survey was deployed using Microsoft Forms. A total of 25 students out of the 175 that participated in the classes responded to the survey giving us a response rate of 14.29%. The students were given the following key instructions: (i) the survey is voluntary, (ii) submissions do not record the students' names, and (iii) the results will be published as part of the broader discussion on these issues.

The students were surveyed with a questionnaire with seven questions using a combination of open-ended (O) and closed-ended (C) questions to evaluate the effectiveness of these lessons. The questions are presented below and were based on LORI, the Learning Object Review Instrument (Gordillo, *et al*., 2020) with some help on the phraseology from Oppenheim (2000):

1. (C) How interesting would you rate the ethics classes?
2. (O) In what way(s) were the ethics classes interesting (if they were)?
3. (C) Did the ethics classes encourage you to look at computer issues from multiple perspectives?
4. (O) In what way(s) did the ethics classes encourage you to look at issues from multiple perspectives (if they did)?
5. (O) What, for you, were the three key takeaways from the ethics lessons?
6. (O) In terms of the content of this lesson, how did you find it?
7. (O) In terms of the format of this lesson, how did you find it?

Overall, the students felt these case studies, including subsequent analysis and discussion, helped them to the consequences of different programming tasks on the end-users and. Furthermore, it also got them to question whether A.I. should be used in situations where people's lives can be affected by the outcome. They also felt these sessions on ethics gave them a new perspective on programming, and they enjoyed discussing the different topics with their classmates and seeing their different viewpoints. They found the topics themselves interesting, authentic, and relevant to their lives and future careers.

4. THE SECOND SURVEY

In this follow-up study six months later, the students were asked to recall the content they could remember from the lessons as well as the key ideas they learned from those case studies. A total of 16 students out of the 175 that participated in the classes responded to the survey giving us a response rate of 9.14%. The students were given the following key instructions: (i) the survey is voluntary, (ii)

submissions do not record the students' names, and (iii) the results will be published as part of the broader discussion on these issues. The questions were as follows:

1. You covered four (4) case studies about ethics in your programming module, what can you remember about any of them? (One sentence about each of them, and if you can't remember any, just say "I can't remember the details of any").
2. Do you think the case studies have made you think about programming and software development in a different way than you had before? If so, how?
3. Have you encountered any ethical scenarios (of any kind) in working with computers (or your phone) since doing these case studies?
4. What are three good pieces of advice you think every person who works in the IT and software field should be aware of, in terms of ethics?

The outcomes were as expected, i.e., some students have almost no recollection of the case studies but remembered general ethical themes, whereas others had detailed recollection of both the content of the classes and the key ethical issues discussed.

5. CONCLUSIONS

The fact that the students had little recall of the detail of the ethics case studies six months after the classes were taught supports Grosz's notion of a "distributed pedagogy", where ethics needs be infused throughout the Computer Science curricula to remind students of some key ethical ideas, as well as to give the students a better understanding of the ethical impacts and possible harmful effects of the technologies they implement (Grosz *et al*. 2019).

ACKNOWLEDGMENTS

The authors of this paper and the participants of the Ethics4EU project gratefully acknowledge the support of the Erasmus+ programme of the European Union. The European Commission's support for the production of this publication does not constitute an endorsement of the contents, which reflect the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

REFERENCES

Brey, P. A. E. (2012). Anticipatory ethics for emerging technologies. *NanoEthics*, 6(1), 1–13.

Canca, C. (2022) "Did You Find It on the Internet? Ethical Complexities of Search Engine Rankings" in *Perspectives on Digital Humanism*, pp. 135-144.

Flanagan, M., Howe, D. C., & Nissenbaum, H. (2008). Embodying values in technology: Theory and practice. *Information Technology and Moral Philosophy*, 322.

Gordillo, A., López-Fernández, D., Verbert, K. (2020). Examining the Usefulness of Quality Scores for Generating Learning Object Recommendations in Repositories of Open Educational Resources. *Applied Sciences*, 10(13), p. 4638.

Gotterbarn, D. (1991). Computer ethics: Responsibility regained. *National Forum*, 71(3), 26.

Grosz, B.J., D.G. Grant, K. Vredenburgh, J. Behrends, L. Hu, A. Simmons, and J. Waldo. 2019. Embedded EthiCS: Integrating ethics across CS education. *Communications of the ACM* 62 (8): 54–61. https:// doi. org/ 10. 1145/ 33307 94.

Hillman, N.L. (2019). "The use of artificial intelligence in gauging the risk of recidivism." *The Judges' Journal*, 58(1), pp. 36-39.

Johnson, D. (1985). *Computer Ethics*. Englewood Cliffs (NJ).

Kelly, A. (2021) "A tale of two algorithms: The appeal and repeal of calculated grades systems in England and Ireland in 2020." British Educational Research Journal, 47(3), pp. 725-741.

Lin, P. (2016) "Why ethics matters for autonomous cars" in *Autonomous Driving*. Springer, Berlin, Heidelberg, pp. 69-85.

Ling, L. and Ling, P. (eds)., 2016. Methods and Paradigms in Education Research. *IGI Global*.

Moor, J. H. (1985). What is computer ethics? *Metaphilosophy*, 16(4), 266–275.

Oppenheim, A. N. (2000). *Questionnaire design, interviewing and attitude measurement*. Bloomsbury Publishing.

O'Sullivan, D., Gordon, D. (2020) "Check Your Tech – Considering the Provenance of Data Used to Build Digital Products and Services: Case Studies and an Ethical CheckSheet", *IFIP WG 9.4 European Conference on the Social Implications of Computers in Developing Countries*, 10th–11th June 2020, Salford, UK.