DEVELOPING AN EDUCATIONAL BRICK FOR DIGITAL ETHICS - A CASE STUDY-DRIVEN APPROACH

J. Paul Gibson, Yael Jacob, Damian Gordon, Dympna O'Sullivan

Institut Mines-Télécom (France), Institut Mines-Télécom (France),
Technological University of Dublin (Ireland), Technological University of Dublin (Ireland)

Paul.Gibson@Telecom-SudParis.eu; Yael.Jacob@Telecom-SudParis.eu;
Damian.X.Gordon@TUDublin.ie; Dympna.OSullivan@TUDublin.ie

ABSTRACT

In this paper we present the concept of re-usable educational bricks for the teaching of digital ethics. After describing the motivation behind the concept, we provide an overview of a standard template that can be used in the design of such a brick. We then briefly review the bricks that are at different stages of development, evaluation, and deployment, following this template. Finally, we conclude with a more detailed review of the development of a brick based on a case study which examines the use of "electronic pills" (e-pills) in the health industry. This case study falls within the computing topic of the Internet-of-things (IoT), and focuses on the ethical issues related to security and privacy.

INTRODUCTION

The work reported in this article is part of the Ethics4EU Erasmus+ Project (more details can be found at http://ethics4eu.eu) -

> The Ethics4EU Project is an Erasmus+ transnational project that will explore issues around teaching ethics in Computer Science. Ethics4EU will develop new curricula, best practices and learning resources for digital ethics for computer science students. It follows a 'train the trainer' model for up-skilling computer science lecturers across Europe.

The project objectives and deliverables are as follows: (i) a research report on *European values in Ethical technology*[1]; (ii) a research report on the *State of the Art of Teaching Ethics in Computer Science programmes*[2]; (iii) a comprehensive curriculum for teaching ethics in Computer Science; (iv) an open access online learning resources database of teaching and assessment strategies for teaching ethics in computer science; (v) an instructor guide to aid the delivery of material from the online resources database; and (vi) an online community of practice to facilitate discussion and experiences in delivering computer science ethics which will complement the online resource database and instructor guide.

The development of the educational bricks is based on the results from research reports (i) and (ii). It is the main contribution to objective (iii) and provides the material for the construction of the online teacher support platform, including a database (iv), instructor guide (v) and community of practice (vi).

[1] http://ethics4eu.eu/european-values-for-ethics-in-technology-research-report/

[2] http://ethics4eu.eu/outcomes/existing-competencies-in-the-teaching-of-ethics-in-computer-science-faculties-research-report/

The remainder of the paper is structured as follows. Section 2 provides the background and motivation for the work. Section 3 introduces the standard template for bricks which is based on a 4-dimensional classification method. Section 4 provides a general overview of the current state of brick development. Section 5 provides a more detailed description of one of these bricks – the *e-pills* case study. Section 6 concludes with a summary of the work that has been completed, and the work that has yet to be done.

MOTIVATION

The importance of well-integrating ethical aspects into computing programmes and modules/courses, as highlighted by Grosz et al. (2019) is well-established; and we are inspired by the research of Chuck Huff and C Dianne Martin (1995) which places emphasis on empathy, and students imagining the consequences of their own work and actions. Furthermore, we wish to encourage a more multi-disciplinary approach to teaching digital ethics as discussed in A.H. McGowan (2012). Our long-term goal is to provide a central repository (platform) of useful re-usable/adaptable education bricks for the teaching of digital ethics, following an "open" model – as proposed by Iiyoshi, Toru, and M. S. V. Kumar (2010) - such as seen with the creative commons approach. This platform will manage teaching material following good software engineering practices - as outlined in J. Paul Gibson, and Jean-Luc Raffy (2011) – for improved maintainability and sustainability.

In order to demonstrate the viability of such a platform we are currently developing a small set of six example bricks. The concept of an educational brick marries closely with that of learning objects, which Wiley (2000) defines as "small (relative to the size of an entire course) instructional components that can be reused a number of times in different learning contexts". It also fits well with the concept of 'distributed pedagogy' as used by Grosz et al. (2019). We hope that with this small set we can generate enough of a critical mass of academic users in order to build and maintain the repository.

EDUCATIONAL BRICKS FOR DIGITAL ETHICS – A STANDARD TEMPLATE

The need for a standard template is vitally important, particularly given the fact that the development of the bricks is being undertaken transnationally, with different bricks being drafted in different countries, and subsequently being reviewed and redrafted in other countries (including in the project partner organisation countries of France, Ireland, Italy, Sweden, and Switzerland).

The repository is currently under development, and the main requirements are for it to provide a rich set of features for searching for bricks, adding bricks, adapting/evolving different versions of bricks, composing bricks, etc. In figure 1, below, we see the brick that is found when we search our prototype system using **HCI** as a keyword: it is concerned with dark patterns in user interface design.

To help standardise and regularise the format and content of the bricks, the template includes two main sections: classification for searching purposes, and pedagogic issues for administrative purposes. Each brick can be associated with one or more case studies, and these are a key part of the classification. With respect to classification, we have four dimensions: ethical issues, academic domains, application domains and interdisciplinarity. Bricks, and case studies, may belong to multiple classes within each dimension.

We note that the template, and associated classifications, is not fixed. As we add more bricks, we expect the template to evolve as we validate its utility for meeting the requirements of the teacher support platform. We note that the platform prototype also permits users to provide feedback on bricks through a comment/chat functionality.

Figure 1. The ETHICS4EU Teacher Support Service Front-End.

Source: http://ethics4eu.eu/brick/dark-pattern-lesson/

Classification of Ethical Issues

The classification of ethical issues in the bricks can be done at two levels of granularity. Firstly, a high-level classification is based on the categories of interest identified in deliverable (ii) of the ETHICS4EU project: a) Origins of Digital Ethics, b) Digital Ethics Values, c) Data Ethics, d) AI Ethics, e) Ethics for Pervasive Computing, f) Ethics for Social Media, g) Governance and Legal Issues, h) Professional Ethics. Secondly, we provide a finer-grained classification based on the identification of ethics keywords in the case studies that we have incorporated in our teaching. This classification will expand and evolve as we add more educational bricks and studies. The current keyword list includes: AI super-intelligence; autonomy; bias, fairness, and transparency; discrimination; intellectual property; privacy and data protection; professionalism; safety and security; cyber-criminality and hacking; society, government, democracy, and environment.

This is the part of our classification scheme that requires more research, including the participation of a community of digital ethics teachers.

Classification of the Academic Domain

For the educational case studies, there is a requirement to match each study with educational requirements. The simplest way to do this is to list the knowledge areas (or skills) which would benefit from students interacting with the case study. The classification of knowledge areas is a complex task, and so we recommend using already developed taxonomies implicit in "bodies of knowledge" or "recommended curricula". The following list of four classification schemes are the most commonly used in Europe (and around the world), and should match with how most educational establishments classify educational content within the domain of "computing".

1. ACM Computing Curriculum[3]

[3] https://www.acm.org/education/curricula-recommendations

2. IEEE Curriculum (SWEBOK)[4]
3. European Research Council's Peer Evaluation (PE6) panel classifications of CS[5]
4. e-skills: The European Foundational ICT Body of Knowledge[6]

Classification of the Application Domain

Outside the academic domain, each brick case study must be classified under one or more application domains. As for the academic domain classification, we recommend using an existing scheme or standard. The three which we have found most useful are:

1. Global Industry Classification Standard (GICS)[7]
2. Industry Classification Benchmark (ICB)[8]
3. ISO Standards[9]

Classification of Interdisciplinarity

The ETHICS4EU approach encourages interdisciplinary teaching. For each case study brick it is very likely that academic disciplines other than computing could be involved in the teaching. For our initial set of brick developments, we have already identified potential for collaboration with the following academic disciplines - biology, physics, electronics, maths, psychology, history, law, medicine, philosophy, and engineering. Including an interdisciplinary classification explicitly acknowledges the opportunity for collaboration with other academic departments and colleagues.

Pedagogic Issues

With respect to pedagogic issues, we have five subsections, which correspond to the type of information that most high-level institutes record with respect to courses, modules and programs that they teach. A sixth is added to explicitly link to other bricks in our repository.

1. Academic Load
2. Pre-requisites
3. Learning Objectives (Ethical, Computing and Transverse)
4. Teaching and Evaluation Approach(es)
5. Support Material (For Teachers and Students)
6. Links to Other Bricks.

[4] https://www.computer.org/education/bodies-of-knowledge/software-engineering

[5] https://erc.europa.eu/sites/default/files/document/file/erc%20peer%20review%20evaluation%20panels.pdf

[6] http://ictprofessionalism.eu/wp-content/uploads/EU-Foundational-ICT-Body-of-Knowledge_Brochure_final.pdf

[7] https://www.spglobal.com/marketintelligence/en/documents/112727-gics-mapbook_2018_v3_letter_digitalspreads.pdf

[8] https://www.ftserussell.com/data/industry-classification-benchmark-icb

[9] https://www.iso.org/standards-catalogue/browse-by-ics.html

A SELECTION OF BRICKS – DEVELOPMENT, EVALUATION & DEPLOYMENT

As part of the ETHICS4EU project, the following bricks are at various stages of development and deployment.

1. *Foundations of Digital Ethics* is a brick which was developed as a pre-requisite to all other digital ethics bricks. Any student who has not had an introduction to ethics as part of their previous education experience will be required to follow the foundations brick. The brick has been developed, and validated by digital ethics experts, but it has not yet been deployed.
2. *Smart Pills* – this brick is detailed in section 5 of this paper. It has been developed and validated, and is due for deployment in May/June 2021.
3. *Software Certification, Accreditation and Testing* - professional ethics for software engineering. This brick is concerned with professional ethics in the domain of transport (aerospace and automobile) and is concerned with testing, certification and accreditation of software systems. The brick incorporates two main case studies – the Volkswagen emissions scandal, and the Boeing 747-Max crashes – which have attracted much media attention in recent years. These studies address the issue of the need for professionalism, and professional ethics, in the development of complex software. The brick is currently waiting for validation.
4. *Introduction to Programming - Algorithmic (AI) Bias* – is a brick for beginners to computing and computer programming. It illustrates that even the simplest algorithms can have bias (not just those based on complex AI and Machine Learning approaches). A central case-study is the use of algorithms for evaluating and assessing students during the COVID-19 pandemic. This brick has been developed, validated and deployed. Currently the teachers are analysing the feedback from the students.
5. *HCI-UX - Dark arts of interface design* – this brick collects a number of case studies from a range of different application domains to illustrate that dark patterns are ubiquitous. The brick examines the sometimes fuzzy boundary between unethical and illegal behaviour with regards to whether the use of such patterns should be considered a criminal activity. This brick has been developed, validated and deployed. Currently the teachers are analysing the feedback from the students.
6. *Autonomous Vehicles - more than just a trolley problem*. This brick is concerned with students' perceptions of autonomous cars, and whether the well-known trolley problem is a good way of teaching about the main ethical issues. The brick has been developed, validated and deployed. Initial analysis of the student feedback is very positive with respect to motivating students to be more aware of and concerned about digital ethics.

The next steps are to expand the community of bricks users. Two complementary approaches are being developed – (i) encourage re-use and evolution of existing bricks, and (ii) addition of new bricks. Already, there has been interest in – and initial development of – bricks which examine the following issues - Student Exam Surveillance and Proctoring, Facial Recognition, Public Surveillance, Tracking, Social Media and Fake news, Professionalism within Teaching and Research Ethics (conflicts of interest, publication practices), and Environmental Ethics - Cloud, AI, NFTs, Crypto-currencies and their impact on the planet.

A MORE DETAILED LOOK AT A BRICK – "E-PILLS"

One of the first bricks developed is one looking at "e- pills" (also known as "smart pills" or "robot pills" or "intelligent-pills"). These are a combination of a drug and a device, which can be described as "an oral tablet that incorporates some type of medical device, such as a microchip, that, for example, controls the release of the active pharmaceutical ingredient after ingestion" (Avery and Liu, 2011). This educational brick is aimed at 3rd/4th year engineering students who have chosen to specialise in information system management and development. As such, they participate in a module concerned with the architecture of complex systems, and apply their learning to developing a prototype system with a real industrial client, as part of a significant team project. In recent years, many of the team projects have incorporated technologies from the Internet-of-Things (IoT). Furthermore, the system requirements have become more and more demanding with respect to data protection and privacy (related to the GDPR in Europe). Finally, the students are becoming increasingly aware of the problem of such systems malfunctioning and the impact on the users.

As part of this module, the students are introduced to published research on general digital ethics issues - Ann Cavoukian et al. (2009), Gauthier Chassang (2017), Nancy Leveson (2020). They are also introduced to ethical issues through mainstream media reports on a wide range of technologies in different application domains. One of these studies is concerned with "smart pills" - Buffy Gorrilla (2017), Sandy Wash (2017). The students are then asked to research the main issues, and are provided with references to general papers on IoT and ethics - Ahmed AboBakr and Marianne A. Azer (2017), Josephina Antoniou and Andreas Andreou (2019) - and specific papers on medical ethical issues - Vinton G. Cerf. (2020), Kobi Leins et al. (2020), Brent Mittelstadt (2017), Julie Myers et al. (2008), Lily Hay Newman (2020), Ziad Obermeyer et al. (2019), Mark Stone (2019), and Daniel Wood et al. (2017).

Through discussion with teaching colleagues and students, there was general agreement that the "intelligent pills" provided an excellent case study with which to develop an educational brick on digital ethics. After playing around with various teaching ideas, the design of the brick was specified using the standard template, as follows. The student workload would be 9 hours contact time + 9 hours independent work. The pre-requisites are foundational knowledge of software engineering and networked/distributed system architectures. The computing learning objectives are: how to read documentation of IOT devices and evaluate whether there is coherency between natural language descriptions, formal technical specifications and the hardware. The ethical learning objectives are: consider who is responsible for the privacy of the sensor data; and the implications of the sensor being faulty/buggy. The transverse learning objectives are: communication skills and interaction with the media. The teaching domains are software engineering, architecture and IoT. The application domain is health. The ethical issues are security and privacy of data. The interdisciplinarity is with journalism and biology. The delivery mechanism/teaching approach is based upon students being involved in a debate with a journalist concerning whether the technical and ethical issues have been well-addressed in the general media. This will involve role-playing, following the advice from Diana Adela Martin et al. (2019). The evaluation is indirect – the students are evaluated through their project work, and one of the criteria is whether they have adequately considered the ethical issues. (The brick is currently being evaluated and refined, for first deployment at the end of the first semester of 2021.)

In figure 2, below, we see the header of the web site specific to the smart-pills brick. In this case the lecturer wished to provide their own front-end for access to the teaching material rather use the default interface provided by the platform. The brick is also included in the platform and will be found using any of the classification keywords; the platform then links to this autonomous web site. We chose not to force teachers to use the platform default template and encourage them to link their teaching material in whatever way is easiest for them. Currently this requires the platform

administration to classify the material by hand, but it is hoped to add automation to support this task. We also note that the web page uses scripts for navigation that are intended to link back to the platform. We hope to do this in the near future and provide a library of similar scripts for platform users.

Figure 2. The E-pills Web page

Source: http://jpaulgibson.synology.me/ETHICS4EU-Brick-SmartPills-TeacherWebSite/index.html

The readers are encouraged to visit the web site for this brick in order to see the different types of material that are provided to the teachers – teaching method support, scientific publications, books, journalistic articles from the popular press, social media posts, video and audio file links, etc. This material is being continually updated.

CONCLUSIONS

This paper has reported in the development of digital ethics educational bricks. This is work in progress, but initial results are very encouraging. We have reviewed the six initial bricks that have been developed, and provided more detail on the brick concerned with e-pills. Much more work is planned for the classification models – we are aware that the ethical classification is just an initial approach in order to quickly facilitate the construction of the teaching platform. We also acknowledge that we need to more formally specify the requirements of the platform in order to aid us in the construction of a community of digital ethics teacher resources. Finally, once the deployment of all initial six bricks has taken place, we intend to carry out an extensive evaluation and share the results.

ACKNOWLEDGEMENTS

The authors of this paper and the participants of the Ethics4EU project (http://ethics4eu.eu) 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.

KEYWORDS: Digital Ethics Education, Case Studies, Internet-of-Things, Smart-Pills, Security, Privacy.

REFERENCES

AboBakr, Ahmed and Marianne A. Azer (2017). *IoT ethics challenges and legal issues*. International Conference on Computer Engineering and Systems (ICCES), pages 233–237. IEEE.

Antoniou, Josephina and Andreas Andreou (2019). Case study: The internet of things and ethics. *The Orbit Journal*, 2(2).

Cavoukian, Ann et al. (2009) *Privacy by design: The 7 foundational principles*. Information and Privacy Commissioner of Ontario, Canada, 5.

Avery, M. and Liu, D. (2011). Bringing smart pills to market: FDA regulation of ingestible drug/device combination products. *Food and Drug Law Journal*, 66(3): 329-352.

Cerf, Vinton G. (2020) *On the internet of medical things.* Communications of the ACM, 63(8):5, July 2020.

Chassang, Gauthier (2017). The impact of the EU general data protection regulation on scientific research. *ecancermedicalscience,* 11.

Gibson, J. Paul and Jean-Luc Raffy (2011). "A “Future-Proof” Postgraduate Software Engineering Programme: Maintainability Issues." *The Sixth International Conference on Software Engineering Advances* (ICSEA 11), Barcelona, Spain (October 2011).

Gorrilla, Buffy (2017) Gut feeling: the swallowable gut sensor that could replace a colonoscopy, *The Sydney Morning Herald.* January 20, 2017. Retrieved from https://web.archive.org/web/20201021122724/https://www.smh.com.au/national/gut-feeling-the-swallowable-gut-sensor-that-could-replace-a-colonoscopy-20170118-gttout.html

Grosz, B.J., Grant, D.G., Vredenburgh, K., Behrends, J., Hu, L., Simmons, A., & Waldo, J. (2019). Embedded EthiCS: integrating ethics across CS education. *Communications of the ACM*, 62(8): 54-61.

Huff, Chuck and C Dianne Martin (1995). Computing consequences: a framework for teaching ethical computing. *Communications of the ACM*, 38(12): 75-84.

Leins, Kobi, Chris Culnane, and Benjamin IP Rubinstein (2020). Tracking, tracing, trust: contem- plating mitigating the impact of covid-19 through technological interventions. *The Medical Journal of Australia*: p. 1.

Leveson, Nancy (2020). *Are you sure your software will not kill anyone*? *Communications of the ACM*, 63(2): 25–28.

Iiyoshi, Toru, and M. S. V. Kumar (2010). *Opening up education: The collective advancement of education through open technology, open content, and open knowledge*. The MIT Press.

Martin, Diana Adela, Eddie Conlon, and Brian Bowe (2019). The role of role-play in student aware- ness of the social dimension of the engineering profession. *European Journal of Engineering Education*, 44(6): 882-905.

McGowan, A.H. (2012). Teaching science and ethics to undergraduates: A multidisciplinary approach. *Science and Engineering Ethics*, pp. 1-9.

Mittelstadt, Brent (2017). Ethics of the health-related internet of things: a narrative review. *Journal of Ethics and Information Technology*, 19(3): 157–175.

Myers, Julie, Thomas R Frieden, Kamal M Bherwani, and Kelly J Henning (2008). Ethics in public health research: privacy and public health at risk: public health confidentiality in the digital age. *American Journal of public health*, 98(5):793–801.

Newman, Lily Hay (2020) Bluetooth-Related Flaws Threaten Dozens of Medical, *Wired*, 20 Feb. Retrieved from https://web.archive.org/web/20201021081330/https://securityintelligence.com/articles/the-potential-and-perils-of-the-iot-in-healthcare/

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464):447–453.

Stone, Mark (2019) The Potential and Perils of the IoT in Healthcare, *Security Intelligence*, November 21. Retrieved from https://securityintelligence.com/articles/the-potential-and-perils-of-the-iot-in-healthcare/

Wash, Sandy (2017) *FDA approves pill with sensor that digitally tracks if patients have ingested their medication*, US-FDA Press release, November 2017. Retrieved from https://www.fda.gov/news-events/press-announcements/fda-approves-pill-sensor-digitally-tracks-if-patients-have-ingested-their-medication

Wiley, D.A. (2000) Connecting learning objects to instructional design theory: A definition, a metaphor, and a taxonomy. *The instructional use of learning objects*, 2830(435), pp. 1-35.

Wood, Daniel, Noah Apthorpe, and Nick Feamster (2017). *Cleartext data transmissions in consumer IoT medical devices*. In Proceedings of the 2017 Workshop on Internet of Things Security and Privacy, pp. 7-12.