

Perspectives on computing ethics: a multi-stakeholder analysis

Damian Gordon and Ioannis Stavarakakis

*Applied Social Computing Network (ASCNet) Research Group,
School of Computer Science, Technological University Dublin – Dublin City Center
Campus, Dublin, Ireland*

J. Paul Gibson

*Institut Polytechnique de Paris, Samovar Lab, TELECOM SudParis,
Évry, France, and*

Brendan Tierney, Anna Becevel, Andrea Curley, Michael Collins,

William O'Mahony and Dympna O'Sullivan

*Applied Social Computing Network (ASCNet) Research Group,
School of Computer Science, Technological University Dublin – Dublin City Center
Campus, Dublin, Ireland*

Abstract

Purpose – Computing ethics represents a long established, yet rapidly evolving, discipline that grows in complexity and scope on a near-daily basis. Therefore, to help understand some of that scope it is essential to incorporate a range of perspectives, from a range of stakeholders, on current and emerging ethical challenges associated with computer technology. This study aims to achieve this by using, a three-pronged, stakeholder analysis of Computer Science academics, ICT industry professionals, and citizen groups was undertaken to explore what they consider to be crucial computing ethics concerns. The overlap between these stakeholder groups are explored, as well as whether their concerns are reflected in the existing literature.

Design/methodology/approach – Data collection was performed using focus groups, and the data was analysed using a thematic analysis. The data was also analysed to determine if there were overlaps between the literature and the stakeholders' concerns and attitudes towards computing ethics.

Findings – The results of the focus group analysis show a mixture of overlapping concerns between the different groups, as well as some concerns that are unique to each of the specific groups. All groups stressed the importance of data as a key topic in computing ethics. This includes concerns around the accuracy,

© Damian Gordon, Ioannis Stavarakakis, J. Paul Gibson, Brendan Tierney, Anna Becevel, Andrea Curley, Michael Collins, William O'Mahony and Dympna O'Sullivan. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) license. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this license may be seen at <http://creativecommons.org/licences/by/4.0/legalcode>

Disclaimer: All authors were involved in the data collection process, reviewing the content of the paper and contributing to the authorship.

Funding disclaimer: This paper is part of the Ethics4EU project which is Co-funded by the Erasmus+ Programme of the European Union under Grant Agreement No 2019-1-IE02-KA203-000665.

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.

Conflicts of interest/competing interests: The authors have no conflicts of interest to declare that are relevant to the content of this article.



completeness and representativeness of data sets used to develop computing applications. Academics were concerned with the best ways to teach computing ethics to university students. Industry professionals believed that a lack of diversity in software teams resulted in important questions not being asked during design and development. Citizens discussed at length the negative and unexpected impacts of social media applications. These are all topics that have gained broad coverage in the literature.

Social implications – In recent years, the impact of ICT on society and the environment at large has grown tremendously. From this fast-paced growth, a myriad of ethical concerns have arisen. The analysis aims to shed light on what a diverse group of stakeholders consider the most important social impacts of technology and whether these concerns are reflected in the literature on computing ethics. The outcomes of this analysis will form the basis for new teaching content that will be developed in future to help illuminate and address these concerns.

Originality/value – The multi-stakeholder analysis provides individual and differing perspectives on the issues related to the rapidly evolving discipline of computing ethics.

Keywords Social impact of ICT, Information communication technology (ICT), Computing ethics, Ethical concerns in ICT

Paper type Research paper

1. Introduction

Computers and technological applications are now central to many aspects of life and society, from industry and commerce, government, research, education, medicine, communication to entertainment systems. The last decade has seen rapid technological growth and innovation, with the realities of artificial intelligence (AI) technology coming to fruition. Topics including privacy, algorithmic decision-making, pervasive technology, surveillance applications and automating human intelligence for robotics or autonomous vehicles frequently undergo scrutiny in the media and are increasingly coming into public discourse. These technologies have wide ranging impacts on society where those impacts can be beneficial but may also at times be negative. There is a sense that some technology development and innovation is happening at a more rapid pace than the relevant ethical and moral debates.

The history of computing ethics (or computer ethics) goes hand-in-hand with the history of computers themselves; since the early days of the development of digital computers, pioneering computer scientists, such as Turing, Wiener and Weizenbaum, spoke of the ethical challenges inherent in computer technology (Weizenbaum, 1976; Bynum, 1999, 2000, 2006, 2018), but it was not until 1985 that computing ethics began to emerge as a separate field. 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).

Deborah Johnson's *Computer Ethics* (Johnson, 1985) was the first major book 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. She also notes the application of deficient common moral norms to new and unfamiliar computer-related moral problems and dilemmas.

In James 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 argues that computer technology makes it possible for people to do a vast number of things that it was not 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, and the concept of "value-sensitive computer design" emerged, based on the insight 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 (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). The code has gone through 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 (ACM, 1992). The last update of the ACM code was in 2018.

Professional bodies continue to play a very important role in producing and disseminating ethical and standard guidelines for ICT professionals, for example the IEEE Ethically Aligned Design Guidelines provide guidance for ICT professionals (Shahriari and Shahriari, 2017). It should be noted that in contrast to other professions such as Medicine or Law, which have codes of ethics and possible penalties in place for noncompliance, the ICT profession still lacks of a coherent umbrella ethical framework (Thornley *et al.*, 2018).

In 1996 the “Górniak Hypothesis” predicted that a global ethic theory would emerge over time because of the global nature of the internet. Developments since then appear to be confirming Górniak’s hypothesis and have resulted in the metaphysical information ethics theory of Luciano Floridi (Floridi, 1999, 2014; Floridi and Sanders, 2005). These new theories make explicit the social and global change created by new technologies and call for an intercultural debate on computing ethics in order to critically discuss their impact on society.

In this paper, we present a literature review of recent work on computing ethics to understand the pertinent topics currently under discussion in the literature and the results of a series of focus groups that explored computing ethical concerns with Computer Science academics, ICT industry professionals, and citizens. The research was conducted as part of the Ethics4EU project’s report on European Values for Ethics in Technology (Ethics4EU, 2021). Ethics4EU is an Erasmus + project that aims to develop a repository of open and accessible educational curricula, teaching and assessment resources relating to computing ethics. The rest of this paper is organized as follows. In Section 2, we discuss relevant recent literature on computing ethics. In Section 3, we describe our focus groups sessions. In Section 4, we present a thematic analysis of the data gathered during the focus groups. We conclude with a discussion in Section 5.

2. Literature review

A systematic literature review approach was employed in selecting relevant literature from a number of key areas that represent some notable present-day computing ethics topics and challenges (similar reviews have been undertaken by researchers such as Braunack-Mayer *et al.*, 2020; Saltz *et al.*, 2019; Saltz and Dewar, 2019). These key areas, also highlighted by Kumar *et al.* (2020), focus on the overlap between three key areas in contemporary computing ethics, the areas of data science, AI and pervasive computing (including surveillance and privacy), thus the focus of this literature review is to critically examine those three areas, and to explore the themes that have emerged in each of those domains in the past five years.

2.1 Data ethics

Data ethics is a relatively new branch of computing ethics that studies moral problems related to data management (including generation, recording, curation, processing, dissemination, sharing and use) as well as algorithms (including those using AI, artificial agents, machine

learning and robots) to formulate and support morally good solutions for data (Floridi and Taddeo, 2016). Data has become a key input for driving growth, enabling businesses to differentiate themselves, and maintain a competitive edge. Value is particularly high from the mass collection and aggregation of data, particularly by companies with data-driven business models. However, the use of aggregated data underpins the risks to individuals' privacy at a very fundamental level. Therefore, it is vital to highlight data management frameworks that promote the ethical collection, processing and aggregation of data. Three popular data management frameworks are the Data Management Association's Data Management Body of Knowledge (DM-BOK) (DAMA, 2017), the Zachman Framework (Zachman, 2008) and the Ethical Enterprise Information Management Framework (E2IM) (O'Keefe and Brien, 2018). A broad overview of the concerns to be addressed by data-based businesses are given in Loi *et al.* (2019) who outline the structure and content of a code of ethics for companies engaged in data-based business, i.e. companies whose value propositions strongly depend on using data.

2.2 Artificial intelligence ethics

AI has emerged as one of the central issues in computing ethics. Pertinent issues in AI ethics include transparency; inclusion; responsibility; impartiality; reliability; security and privacy. While many researchers and technology experts are excited by AI's potential, many others are unsettled by it. Authors balance the positive effects of AI (self-driving cars leading to better safety, digital assistants, robots for heavy physical work; and powerful algorithms to gain helpful and important insights from large amounts of data) against the negatives (automation leading to job losses, rising inequalities attributed to AI haves and have nots and threats to privacy) (Helbing *et al.*, 2019).

Algorithmic transparency and fairness are key elements of ethical AI systems (Webb *et al.*, 2019). Ethical AI systems should also work to eliminate bias which can be achieved by a greater understanding of the data used to build systems (Floridi and Taddeo, 2016; Mittelstadt *et al.*, 2016; Müller, 2020). A key element for automated decision-making systems is accountability and auditing (Whittaker *et al.*, 2018). If the system involves machine or more recently, deep learning, it will typically be opaque even to the expert who developed it. Machine and robot ethics are another important subfield of AI ethics. If a robot acts, will it itself be responsible, liable, or accountable for its actions? (EGE, 2018). "The effects of decisions or actions based on AI are often the result of countless interactions among many actors, including designers, developers, users, software, and hardware. [...] With such distributed agency comes distributed responsibility" (Taddeo and Floridi, 2018, p. 751).

2.3 Pervasive computing ethics

The terms "pervasive computing," "ubiquitous computing," "ambient intelligence," and "the Internet of Things" refer to technological visions that share one basic idea: to make computing resources available anytime and anywhere by embedding computational devices in everyday objects, freeing the user from the constraint of interacting with ICT devices explicitly via keyboards and screens.

One of the central tenets of pervasive computing is "Understanding and Changing Behavior," a topic that clearly has significant ethical considerations (Kranzberg, 2019). Related elements include surveillance technologies, effects on privacy and technological paternalism (Hilty, 2015; Macnish, 2017).

The ethics of surveillance considers the moral aspects of how surveillance, including facial recognition technology, is used (Macnish, 2017). One of the core arguments against surveillance is that it poses a threat to privacy and in a world of ubiquitous automatic

identification, the amount of personal data generated and circulated is expected to increase dramatically.

Privacy is an integral part of pervasive computing ethics and is defined as an individual condition of life characterized by exclusion from publicness (Rasmussen *et al.*, 2000). In the context of computing, privacy is usually interpreted as “informational privacy,” which is a state characterized “by controlling whether and how personal data can be gathered, stored, processed or selectively disseminated”. There is a clear conflict between privacy and pervasive computing technologies, particularly those technologies that deal with sensing and storage (Jacobs and Abowd, 2003). The resulting requirement to protect individual privacy against data misuse entered many laws and international agreements under different terms, some of them focusing on the defensive aspect, such as “data protection,” others emphasizing individual autonomy, such as “informational self-determination.

3. Methodology

Data collection was conducted during the Ethics4EU first multiplier event in November 2019 inside the campus of Technological University Dublin (TU Dublin) in Dublin, Ireland. Participants from academia and industry were recruited as part of a convenience sampling approach (Saumure and Given, 2008). TU Dublin’s School of Computer Science has extensive collaborations with the ICT industry. Therefore, readily accessible participants from a list of contacts were invited according to their field of work in a range of ICT organizations (both large and small). Academic participants from a range of academic institutions were identified through their research areas and expressed interest in the topic. Citizen participants were recruited following a snowball sampling method (Morgan, 2008b). Some participants were initially contacted by the researchers and were asked to spread the word to other citizens who were interested in participating in the focus groups.

A focus group approach was used to help identify digital ethical concerns from stakeholder groups. Focus group interviews give researchers the opportunity to acquire simultaneously a variety of perspectives on the same topic (Gibbs, 1997). Zikmund (1997) identifies ten advantages in the use of focus groups for research. Among these are speed, stimulation (participants’ views emerge from the group process), security (because of the homogeneity of the group) and synergy (when the group process brings forward a wide range of information) to name a few.

These advantages were aligned with our study where focus groups were used for the following reasons. Firstly, “Ethics” as a branch of philosophy is inherently about argumentation, debate, discussion and negotiation on things like practical moral issues, the morality of one’s actions and on the definition of morality itself. As there are no black and white answers, communication and argumentation are fundamental properties in ethical reasoning. Therefore, in an attempt to approach it as an “alive”, interactive and ever-evolving process focus groups were considered to be a good methodological tool for capturing this angle. To our knowledge, this has not happened before in the computing ethics literature.

Secondly, our study is exploratory in nature and our aim was to investigate the topic broadly and not so much in-depth (Stokes and Bergin, 2006). The aim is that our focus groups will be supplemented later in the research with more in depth interviews. Thirdly, groups can be less expensive and faster for data collection compared to one-to-one interviews. Fourthly, compared to one-to-one interviews where the researcher can be at risk of biasing the interview (Vyakarnam, 1995), focus groups reduce risk mainly because of the more balanced social dynamic within the group. However, this is dependent on the style of the group moderators and on group homogeneity. Homogeneity (Morgan, 2008a) is based

upon the shared characteristics of the group participants, relevant to the research topic. In this case, for example, one common characteristic was the professional field, e.g. academics were placed in one group and people from industry in another.

Each group were asked to discuss their views on three open questions (see below) and spent approximately 30 min discussing each question. The number of participants in each focus group was kept at 10–12, and this meant that there was one industry group (10 participants), but it necessitated creating two academic groups (11 and 12 participants) and two citizen groups (11 participants per group).

Each focus group began with an introduction to the main goals of the Ethics4EU project. Consent was obtained from each participant with the clear understanding that they have the right to withdraw at any time. The participants were also informed that their privacy would be respected, and the data from this research would be secured in a protected location, with adherence to GDPR. This introduction process took 10–15 min, depending on the number of questions from the individual groups.

Each focus group was assigned a moderator. In our study, the group moderators took a semi-directive approach allowing members of the group to speak freely while encouraging all members to participate for each of the three questions. According to [Morgan \(2008a, p. 354\)](#) “[t]his strategy matches goals that emphasize exploration and discovery”.

The moderator facilitated the introduction, introduced the questions and encouraged all participants to contribute. They also made sure no single person dominated the discussion. They did not attempt to steer the discussion in any direction, rather letting the topics emerge from the participants. In addition, moderators took detailed notes of the discussion using pen-and-paper. Audio recordings were not taken.

Three questions were developed to garner the participants’ specific concerns and insights into computing ethics issues. We were interested in gaining an overview of topics that participants considered important and to analyse if these topics were reflected in the computing ethics literature. We did not seed any ideas with the participants; rather we allowed an open discussion and topics to emerge. As part of the Ethics4EU project, our aim is to developing learning materials for computing ethics and as such we were also interested in participants ideas regarding teaching and training with regard to computing ethics. The three questions were as follows:

- Q1. What ethical concerns do you have about new technologies?
- Q2. What skills or training should people have to protect themselves in the online world?
- Q3. What ethical training should be given from persons designing and developing technology, and who do you think should give that training?

One reported downside of focus groups is the tendency for some participants to conform to group consensus ([Stokes and Bergin, 2006](#)). Also focus groups cannot delve deep into each participant’s range of opinions and beliefs for any given topic. However, these disadvantages have been acknowledged from the start and were not deemed to negatively interfere with the scope of the current study.

3.1 Participant demographics

3.1.1 Industry participants. The majority of the participants (80 %) were aged 30–49, and 20% were aged 50–69. Half of the participants (50%) were female and 40% were male with 10% preferring not to say. Thirty percent of industry participants had a bachelor’s degree, 60% had a Master’s degree, and 10% had a PhD.

3.1.2 *Academic participants.* A majority of the participants (65 %) were aged 30–49, 13% were aged 18–29, and 22% were aged 50–69. A total of 48 % of the participants were female and 43% were male, with 9% preferring not to say. A total of 9% had a bachelor’s degree, 39% had a master’s degree and 52% had a PhD.

3.1.3 *Citizen participants.* A majority of the citizens (82 %) were aged 30–49, 14% were aged 18–29 and 4% were aged between 50–69. A total of 64 % of the participants were female and 32 were male with 4% preferring not to say. A total of 18% had second level education, 50% had a bachelor’s degree and 32% had a master’s degree. The professions of the citizen participants are shown in [Table 1](#). If there was more than one participant in a given profession, the number is shown in brackets after the profession. A total of 17 professions are represented among the participants.

4. Results

Participants’ responses to the three questions were analysed using a thematic analysis, which is an approach for identifying themes or patterns of meanings with qualitative data. The key step in thematic analysis is coding the data, which involves attaching labels (or codes) to phrases or sentences of analytic interest. In this research a modified version of the coding process based on [Gorden \(1992\)](#) was followed:

- *Define the coding categories:* When the focus groups were completed, the researchers familiarized themselves with the data by reading the transcripts many times. The researchers looked for patterns and themes across the data. Initially a *colour coding* approach was used to identify the main themes emerging from the transcripts. This involves highlighting different parts of the transcripts in different colours to represent initial themes. This gave the researchers the ability to look at emerging themes “at a glance”, and to explore the balance of text that relates to each theme.
- *Assign code labels to the categories:* From this first step, a preliminary tentative set of text codes were created to describe the computing ethics topics emerging from the transcripts. These codes replaced the coloured text. Examples of the initial codes included: *digital-literacy, where-law-overlaps-ethics, older-people-concerns*
- *Classify relevant information into the categories:* Following this initial process, the transcripts were re-read and the initial codes were attached to all relevant text. It was found that in some cases these codes were too general, e.g. one of the early

Managers	Healthcare	Education	Others
<ul style="list-style-type: none"> • Category Management • IT Manager (2) • Marketing Manager • Operations Manager • Project Manager (2) 	<ul style="list-style-type: none"> • Dentist • Carer • General Practitioner (2) • Naturopath 	<ul style="list-style-type: none"> • Primary School Teacher (2) • Researcher • Student 	<ul style="list-style-type: none"> • Freelance writer • Quality Engineer • Recruitment Services • Sales (2) • Secretary

Table 1.
Professions of citizen participants

themes was: *data-ethics*, which was later deemed to be too general; and in other cases the code were too specific, e.g. *ACM-professional-code-of-ethics*.

- *Refining the codes*: Following the identification of codes that were not fully suitable, those that were too general were further refined, e.g. *data-ethics* became *data-ethics-privacy*, *data-ethics-reliability*, *data-ethics-retention*, and *data-ethics-misuse*, and those codes that were too specific were merged, e.g. *ACM-professional-code-of-ethics*, *employee-responsibilities* and *organizational-specific-guidelines* became *professional-ethics*. Some other codes were renamed, e.g. *relevant-European-laws* became *importance-of-GDPR*. This was an iterative process and took place over a period of three weeks to complete.
- *Test the reliability of the coding*: The reliability of the coding was tested by asking an independent reviewer to code one of the transcripts without having access to our coding process (this is called an *independent-coder* method). There was a strong overlap between the two coding processes, thereby validating the approach.

The final codes are the key themes of the transcripts, which are described in the subsections below. Each theme is highlighted in bold. Responses from the different groups are presented comparatively for each question. Where there were multiple groups of participants, we have combined the responses.

4.1 What ethical concerns do you have about new technologies?

4.1.1 *Industry responses*. There was agreement amongst the participants that one of the major areas of concern was the *on-going automation of activities that were previously undertaken by human beings* (this concern was further exacerbated for the group when the automation is achieved through Machine Learning). Participants discussed what to do with the people when their jobs are being replaced by machines - “we’re developing technology that is taking jobs away from people, and although they can reskill, it is not clear that enough new jobs will be created to replace those that are going to be lost”.

One of the most discussed considerations was the challenge of *bias in automated decision-making systems*, and this issue was looked at from the interrelated perspectives of bias in datasets, and bias in machine learning algorithms. When discussing bias in datasets participants highlighted the potential dangers of using open datasets, which may not have been analysed for completeness, and “may exclude particular populations, for example, those on the margins” and yet the conclusions derived from the analysis of these datasets may be presented as fact. There was also a good deal of discussion of potential historical patterns of bias in datasets (including issues around gender and race), and how to prevent those historical issues from being propagated. Suggestions for solutions included “exploring patterns for bias in data, looking at statistical variances, examining who owns or controls the data and how the datasets were created”. Participants suggested that this type of analysis “should look at composite biases that are more difficult to detect, for example, hiring women over a certain age in employment practices”. It was also suggested that the “GDPR guidelines are useful for exploring bias”.

There were discussions around *environmental considerations in the collection of that data*, particularly in the context of Internet of Things. Another participant mentioned the role of data centers contributing to the environmental impact.

One organisation that was highlighted for their excellence in ethics was German Software Company, SAP SE, who incorporate a great deal of ethics in their graduate training programmes. A representative from SAP described how their training “provided three examples of where there had been ethical compliance breaches, and one example

looked at a senior female colleague working in Israel and Korea and how she was treated differently in different countries.” This discussion led onto a conversation on *regional and cultural differences in ethical standards*.

Participants discussed *IT professional standards and their relation to ethical standards*, in particular if IT professional standards can ensure that ethical standards are adhered to. Some participants felt it was important to note that very often employees do not necessarily have control over their work and their use of data and code, and may not be aware of where their work will be used or how it can be repurposed and that there is no clear guidance in IT professional standards in such cases.

Finally, all participants agreed that students should be educated *about legal frameworks with relevant ethical aspects* including Confidentiality Agreements, Non-Disclosure Agreements and Intellectual Property.

4.1.2 Academia responses. The academic participants began by discussing their concerns about *data and datasets* – one participant remarking “I worry about the amount of personal information that is being kept, I think we are keeping a lot more data than we need.” Both groups felt that a lot of organisations seem to be collecting as much data as possible with no clear purpose, other than feeling there is something of value in the data. They also expressed their concerns about *completeness and representativeness of datasets* (particularly for open source datasets), and the *potential bias that an AI system might embody* in its decision-making if trained on these data sets. This led to a discussion on the ownership of decisions in an online context, for example, if a search engine is suggesting search phrases, and suggesting sites to visit, who is really making the decisions – one participant said “where is the line between me and the application?” Another concern raised was the dangers of *technological or digital colonisation in developing countries* and a lack of control of their data on the part of citizens in the developing world.

Another concern was that data sets might exclude certain people for *privacy* reasons because they might be identifiable due to their specific characteristics, and therefore could be unrepresented in the overall data sets. This conversation about *marginalisation* led to a discussion of how some voice recognition technologies, for example, one participant remarked how voice assistants such as Alexa, might not be as effective for people with speech impediments or strong regional accents.

On the theme of *privacy*, participants expressed concern over the possibility of governments or private organisations combining various data “personal, legal and financial information” about an individual and the impact that might have. One example cited by a participant was the South Korean government limiting the number of (and size of) gambling transactions that any one person can take part in per day when gambling online. Even though there was general agreement that too much gambling is a bad thing, nonetheless the notion of the government being able to restrict an individual’s liberty was considered objectionable, and they felt that education and help is preferable to controlling them. Participants remarked that a combination of control and educational measures have been used by governments to reduce the number of people who smoke.

Both groups highlighted the importance of cybersecurity as being an ethical imperative in the context of the significant amount of personal data being collected on individuals, as well as the numerous high profile data leaks that have occurred.

Both groups discussed research projects and the importance of *ethics for research projects*, particularly when funded by public monies. They discussed the importance of transparency in research projects and how this should be communicated to the public. For example participants discussed how it is important to make clear “whether enticements or incentives were permitted.” They discussed the importance of research communication more

generally and how this is an ethical issue, for example how research findings are presented to the general public who may not be familiar with the general area or the specific details. Finally, both groups agreed on the importance of teaching students about plagiarism, copyright and the honest presentation of results and outlining how such practices are a violation of ethics as it is unfair to take credit for another person's intellectual property or to present misleading results or information.

4.1.3 Citizens responses. Many of the concerns of the citizen participants related to *data*. Some common themes were around *collection and retention of data* and *misuse of data*, in particular how data collected on individuals may be used and misused. Data gathering with the purpose of creating profiles, for example voter profiles with the purpose of influencing democratic elections were mentioned frequently by the participants.

Privacy was another topic discussed by all participants, in particular data appropriation on behalf of privately owned businesses, which do not reward their users for the acquisition of such data and use it to maximize their profits. Others raised concerns about *third parties gaining access to data without the original users consent*. Concerns were raised about applications (e.g. Facebook, Siri) that can listen to conversations through digital devices and use that information for targeted advertising. The balance between commercial gain and social benefit was discussed.

Respondents cited a number of *concerns related to social media* including the *normalisation of unacceptable behaviours*, inappropriate exposure to information and content, *widespread dissemination of misinformation*, the fuelling addictive tendencies and preying on vulnerable individuals. Social media platforms can have negative affect on individuals psyche leading to mental health issues and reduce genuine human interactions and empathy. Particular concerns were raised about the influence of social media on teenagers and young people, for example cyberbullying and a lack of awareness on the part of young people about the longevity of data online. One participant said "people who come into the public eye are likely to have their social media postings from decades earlier where they were perhaps a younger and less-informed person, examined for any failures to be used against them". There were also concerns about misinformation online and the sharing of "fake news" via social media, and the detrimental impact that can have on younger, most impressionable people, one participant noted "As a parent I wonder how safe my teenagers are online, they lack the experience and skills to distinguish reputable sources from fake news".

The participants also felt that *surveillance of individuals, in particular using facial recognition technology* was of great concern, where surveillance companies are potentially storing images and information about people who are unaware they are captured on camera, for example from walking on the street, going shopping or entering other commercial buildings. One participant notes that there are cultural differences in how surveillance technologies are being used "Surveillance is creeping up in prevalence but thus far I believe it is being used for legitimate and good purposes, for solving crimes, preventing crimes, and finding missing persons, in Ireland at least. I wouldn't be quite so sure about other countries and other one-party states."

4.2 What skills or training should people have to protect themselves in the online world?

4.2.1 Industry Responses. The group analysed this question from the point-of-view of their work as software designers and developers. There was a discussion centred around the notion of *consequence scanning*, where designers and developers try to predict the consequences of developing the software they are asked to create. For example, designers and developers should consider what the software should and should not do, the worst

possible negative consequence of the software and how the software would work if repurposed for another system.

Generally, the participants reflected that a lot of employees in organisations do not get to see the “big picture”, and therefore do not have the opportunity to evaluate the ethical implications of the processes that they are involved in. On the other hand, managers who have the bigger picture, but do not know the exact detail of how systems have been developed might also be missing out on some of the ethical implications of the how the work of different designers and developers impact each other from a moral point of view.

Another consideration that was discussed was the *dangers of using off-the-shelf code*, particularly when used by naïve or novice designers and developers, who may not have thought of the complete ethical implications of using that code, or may not have full information on how the off-the-shelf code works, and therefore have no awareness of the potential ethical issues. The conversation moved to the important role that educators must play in exploring these issues.

The participants reflected that there is a need for *more diversity in the IT profession*; importantly as trainers, as designers; as developers, and as testers, so that “they can ask the ethical questions that others don’t think of”.

4.2.2 Academia responses. The groups looked at this question from the point-of-view of teaching students how to design and develop a computer system. The groups, began by discussing considerations that students should have before designing and developing a computer system. These considerations could broadly be characterised as *consequence scanning*. For example, what is the best outcome of this development? What is the worst unintentional outcome that could happen? How would I mitigate the worst outcome if it happened?

There was a general agreement of the importance of “always keeping a human in the loop”, and particularly to ensure that there is consultation with persons with significant domain knowledge as they will understand in what context the technology will be used. There was also agreement that we have to encourage designers and developers to think more reflectively and “consider what the system they are developing is really for, and is really about”. One participant suggested a possible scenario where a developer was asked to create software following a specification, and it became evident that the system as a whole was designed to make the software addictive, even though it was not evident to the individual developers, what should they do?

There was also general agreement that *explainability* is vital in all *automated decision-making* processes. This explainability concept refers to both the ability to understand terminology such as “features” and “weights”, but also that each individual decision that the system takes can be explained. As part of this conversation, participants questioned whether or not it is appropriate to develop systems using partially correct data. Another participant mentioned the use of software libraries like LIME for Python which helps explain how machine learning systems make specific predictions.

The groups agreed that there is a need for designers and developers to be aware of the law as it pertains to them, and *where the law overlaps with ethics*. Although there was general agreement that often ethical principles can be of a higher standard than the law, but the groups wondered if there is one set of ethical principles that should be followed by everyone. Further to this, there are different laws in different countries, and there may even be different ethical standards in different regions that developers should be aware of. The topic of outsourcing was discussed, where systems can be developed in one region but used in another and different ethical standards can apply in the different regions.

The groups also discussed the nature of ethical standards, and wondered where can you find (and find out about) standards and *how ethical standards can be enforced*. They also questioned *whether ethical standards can keep up with the rapid developments of software*. An issue discussed was the impact of unethical behaviour which impacts commercial activities, but more importantly impacts people and consumers. One participant highlighted *opaque Terms and Conditions* that many users sign up to without reading, and agreeing to things they either do not read or do not fully understand.

Another key issue discussed was *accessibility* and the importance of ensuring that as wide a range of people as possible can use the software being developed. One participant commented “Sometimes the client might not be aware of, or concerned with, accessibility consideration, but does that mean the developers shouldn’t consider it?”

Finally there was some discussion on how equipped academics are to teach this type of content, and *what sort of training or teaching content is required by academics* so that they can become confident in teaching this topic. Both groups felt that such content should be publicly available to private and public organisations.

4.2.3 Citizens responses. Participants in this group took a broad view of the question. The issue of the *longevity of digital information* was raised as a concern, particularly social media posts which may be innocuous in the context in which they were created, but could be potentially misunderstood or misrepresented without that context, and could be detrimental in the future. As one participant phrased it: “For all groups they need to be made aware, understand implications of posting information to a world that is never deleted, follows them around forever, for example, years old tweets coming back to haunt people and impact on work opportunities”.

Participants felt that people should *understand how data can be obtained* by others and be taught about their digital security and online platforms including privacy settings. As one participant said: “People are unaware they are ‘the product’ in many cases. I think the phone companies and social media companies need to be much more transparent when people open accounts about ethical issues around obtaining” Another participant gave the example of digital assistants like Siri or Alexa “always listening to conversations in the home and using the information for marketing purposes”. At a more fundamental level *digital literacy* was considered extremely important, one participant remarked “People should know the basics of digital literacy, cookies come to mind. I don’t fully understand these, yet every website asks to allow them be used.”

Both groups agreed that parents need specific training to help them navigate the online world and to understand the implications of having a digital presence. They suggested that training should include using parental controls (including monitoring tools) and other ways of securing devices (in hardware and software), knowing some of the key social media applications, how to deal with cyberbullying, and how to talk to their children “about the positives and negatives of social networks, and ensuring they keep the channels open to enable children to discuss issues or bumps they encounter in their cyber journeys”. Another participant suggested that parents and children “should be taught about risks, to have an awareness of strangers online and false accounts or information”. They also felt it would be extremely helpful to learn about the addictive nature of social media applications and smartphones, and advice on how to limit their children’s use of these technologies. They stated it would be helpful to know what is legal and illegal in terms of sharing and downloading of audio and video files.

Both groups also agreed that *older people need training to help them navigate the online world*, particularly if it could be tailored for their interests, including lifestyle application (health, banking, shopping) and privacy settings on their devices and applications.

Most participants felt that training about scams and fraud would also help, as well as general personal data protection online. All participants felt that some older people may not want to (or be able to) access digital services, and therefore offline services (government services, libraries, postal services) should be maintained for this age group if that is their preference. One participant commented that “the move to online services is regrettable particularly since some older people may not have a laptop or smartphone, may not have an internet connection (their area may not have coverage), or may not be comfortable using online banking applications, and maybe therefore be far more vulnerable to phone scams”.

Participants also mentioned that “voluntary organizations, clubs, societies also need training on the do’s and don’ts of social media and use of members’ data”. They also felt these groups should be trained on how to recognize and report inappropriate content. Finally they felt that “all elements of security awareness from spamming, phishing, identifying secure sites, should be taught, with real life examples”.

4.3 What ethical training should be given from persons designing and developing technology, and who do you think should give that training?

4.3.1 Industry responses. The first topic that the participants discussed in some detail was *the importance of GDPR* – (the General Data Protection Regulation GDPR 2016/679), and the importance of ensuring that students know how to perform data protection impact assessments, be able to handle data securely, and realise the importance of handling sensitive data before graduation. There was also agreement that a work placement during a university course can be very beneficial in terms of learning the importance of computing ethics, and can teach the students lessons that may be more difficult to teach in the classroom.

The discussion then turned to what other skills need to be and participants suggested “how to develop empathy” and “respecting social norms” as fundamental skills. Additionally there was agreement on the importance of giving students that ability to deal with situations where they are asked to do something unethical, including “the tools to ask further questions in situations where there appears to be unethical activities occurring to gain a deeper understanding of the situation”. Participants also noted that another way to help students develop an appreciation of computing ethics is to help them understand the reason why ethics are being breached.

Participants felt that the most effective way to successfully teach ethics would be to *incorporate ethics into existing modules* rather than creating a dedicated separate module, and it was suggested that something as simple as writing a small reflective piece on how data protection or computing ethics applies to a specific module might be a way of raising awareness, and to consider having someone external to the module review the content of these pieces.

Finally everyone agreed that ethics is a broad societal issue, it is up to everyone to help develop their personal, and societal, understanding on ethical standards.

4.3.2 Academic responses. The groups stressed the importance of *communication, teamwork, and most especially a sense of personal responsibility as key skills* that need to be taught as part of computing ethics courses. One participant mentioned the importance for graduates to understand the concepts of “*informed consent and voluntary consent*”. All participants agreed that “confidence is also a vital skill in graduates, including having the confidence to ask difficult questions of colleagues and of management”, as well as having the confidence to speak up when unethical issues arise, and having the confidence to adapt to a changing environment. A related discussion involved how to equip graduates with the

ability to think critically about their work once working in a company, and to appreciate that there may be conflict between ethics and profit.

Participants also felt that it was vitally important that graduates be equipped with a good understanding of data science giving the recent advances in the topic. Participants felt it is important for students “to explore the power and dangers of aggregate data, including a discussion of how to make an individuals’ data private using techniques such as differential privacy”. Another participant commented that students should be equipped with “an understanding of both bias and representativeness of datasets”.

The groups discussed *the relationship between ethics and the law*, and how knowledge of GDPR legislation is very important for all graduates. From there the discussion moved onto the potential conflict between legal issues and ethical issues, and what choices the graduates should make under those circumstances. One participant commented “it is important that every organisation should have a clear code of ethics, and promote that code, and promote ethical thinking in their organisation”.

4.3.3 Citizen responses. The participants felt that the key skills for people designing and developing technology is “sympathy and empathy to think about the people who will be using the technology”, and that it is “important to realise that not everyone is a technology wizard”. They felt that using new technology can be difficult for many end-users and that many of them may not even think about or understand the ethical implications of technology. The participants also noted that even if they fully understand how to use the system “that doesn’t mean I fully understand how it works, so there might be a whole layer of ethical issues that are not visible to me”.

The groups felt that *ethical training about “laws, codes, and policies”* are also very important, and both groups mentioned that the people designing and developing technology must also have the confidence and courage to ask questions of their organisations, and ask questions of themselves to ensure that the highest ethical standards are being adhered to.

The majority of participants thought that Universities should be responsible for the ethical training of persons designing and developing technology. Some did suggest the responsibility lies with employers, and a few thought it should be part of a continuing professional development while others thought it should be an individual’s personal responsibility.

4.3.4 Overlapping concerns among all participants. The final codes from the focus groups data were collated and common themes between participant groups were identified. It is clear that there were specific thematic overlaps among all groups – privacy, consequence scanning and where the law overlaps with ethics. Overlapping themes between academics and industry people are bias in automated decision-making systems, regional and cultural differences in ethical standards and the importance of GDPR. Overlapping themes between academics and citizens were misuse of data, widespread dissemination of misinformation, surveillance of individuals – facial recognition technology and accessibility (Figure 1).

Common themes mentioned by industry professionals and citizens were automation replacing human beings and longevity of digital information. Themes only discussed by one group are also shown and some interesting findings include that academics alone were concerned about completeness and representativeness of datasets, the impact of technology in the developing world, explainability of automated decision making; industry professionals alone discussed environmental considerations in the collection of data, the dangers of using off-the-shelf-code and a lack of diversity in the ICT industry; citizens spoke about third party access to data, digital literacy and how older people may need help to navigate the online world.

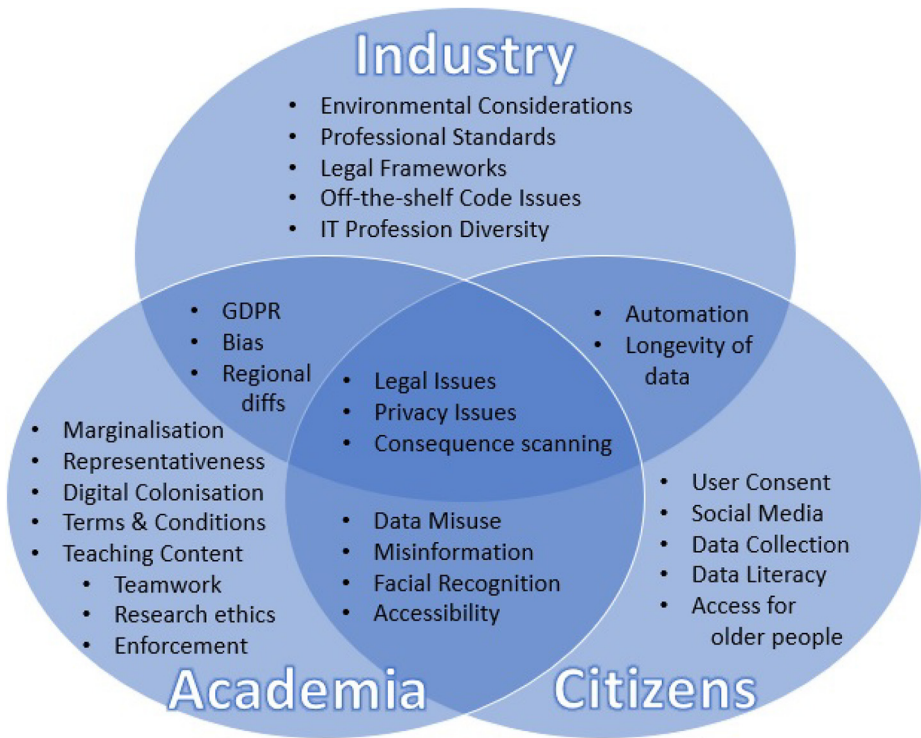


Figure 1.
Overlapping
concerns between
participants from
focus groups

It should be stated that although the sample size was reasonably large it was not investigated *why* some issues were mentioned by one group and not by another. Therefore, we must be tentative in our conclusions and not make any generalisations outside of our sample.

5. Discussion and conclusions

In this paper we have presented a literature of pertinent computing ethics topics from 2014–2019 and the results of a multi-stakeholder analysis where we examined the computing ethics issues that are deemed most pertinent by computer science academics, ICT professionals and citizens. The focus groups showed a combination of overlapping concerns as well as some concerns that are unique to each of the specific groups. All expressed concerns around data, in terms of privacy and data collection and secure storage, bias in datasets and data misuse. All groups also expressed concerns that developers often lack empathy and do not fully understand the user groups they are developing technology for. Academics were concerned about how computing ethics is taught to computer science students, often as a standalone course or module that does not reflect the distributed and interrelated nature of computing ethics concerns that cut across many computer science topics. Industry participants highlighted legal aspects including the importance of GDPR and the legitimate use of data. Citizens expressed a broad range of concerns about social media applications including concerns that social media technology has led to the normalisation of unacceptable behaviours, inappropriate exposure, and preying on

vulnerable individuals. Citizens also highlighted the need for training in an online world, for example, how to deal with cyberbullying and how to identify possible scams and fraud.

The topics discussed by the focus groups overlap well with our findings from the computing ethics literature – concerns around data ethics and automated decision-making systems as well as about privacy and influence of social media were voiced by participants. Participants also discussed topics less well developed in the literature including the environmental impact of computing, the enforcement of ethical standards, the role of personal responsibility in developing technologies and training needs for specific groups.

It is clear from the analysis that there is a broad range of computing ethics concerns and that all stakeholders are considering the dilemmas, pitfalls and solutions. The focus groups considered contemporary topics, many of which have only fully emerged in the last decade. It is very likely that new technologies with a new set of ethical dilemmas will emerge soon. For example, this work was conducted in 2019, prior to the COVID-19 pandemic where much has been written about the ethics of contact tracing technologies and associated privacy concerns.

The goal of this research was to collect key computing ethical issues that are of concern to three stakeholder groups to help develop teaching content for students in computer science programmes. These three stakeholder groups represent the transition that many undergraduate students will undergo; before their programmes they are *Citizens* (knowing little about the detailed mechanics of how computers work, and the ethical issues associated with them), during their programmes they become *Academics* (knowing more about how computers work, and discussing ethical issues from an academic perspective), and when they graduate, they become members of the *Industry* group (Learning how computers work in the professional environment, and which practice ethics issues come to the fore).

The Venn diagram therefore can be seen as a set of themes or motifs that can be incorporated into a computer science programme, that can add substantive ethical content into that programme. The exact sequence of teaching the ethical content will depend on the overall nature of the programme, but as the students transition from novice to expert learners, the depth and complexity of the discussions of ethical issues, that they can have, can continuously grow and evolve.

Researchers such as Moore (2020) advocate that a computer ethics curriculum needs to be dynamic, evolving, and relevant to students' lives and their beliefs, and therefore topics that are of genuine concern to them, and in particular the political nature of computing technologies (including many of the topics highlighted by the stakeholder groups) can be used as a powerful source of motivation in educating future generations of computing students.

As stated earlier, in our experience, the use of focus groups for investigating computing ethics is a novel one but it can prove particularly useful because it provides qualitative insights on the ethical reasoning processes of the interplay between academics, industry and citizens on the moral issues of rapidly evolving digital technologies.

While technological innovation has many positive aspects, we should not be blind to ethical and moral imperatives and strive to develop responsible technology as a first principle. Otherwise, it can prove difficult to reverse the consequences. There is a responsibility on universities to ensure these blind spots do not exist and that students who design and develop technology are taught to consider both the expected and unexpected, and positive and negative consequences of any systems they implement. We will address these concerns as uncovered via the literature and the analysis of the insights from our focus group as part of our research in the Ethics4EU research project. More specifically, the project consortium are developing educational material in the form of lessons that address

topics such as bias in AI, the creation, use, longevity and environmental impact of datasets, programming errors, accessibility, privacy and facial recognition, and ethics of smart devices and pervasive computing. Initial educational materials are available via the Ethics4EU website at <http://ethics4eu.eu/bricks>.

References

- ACM (1992), "ACM code of ethics and professional conduct", Code of Ethics.
- Braunack-Mayer, A.J., Street, J.M., Tooher, R., Feng, X. and Scharling-Gamba, K. (2020), "Student and staff perspectives on the use of big data in the tertiary education sector: a scoping review and reflection on the ethical issues", *Review of Educational Research*, Vol. 90 No. 6, pp. 788-823.
- Brey, P.A.E. (2012), "Anticipatory ethics for emerging technologies", *NanoEthics*, Vol. 6 No. 1, pp. 1-13.
- Bynum, T.W. (1999), "The development of computer ethics as a philosophical field of study", *The Australian Journal of Professional and Applied Ethics*, Vol. 1 No. 1, pp. 1-29.
- Bynum, T.W. (2000), "The foundation of computer ethics", *ACM SIGCAS Computers and Society*, Vol. 30 No. 2, pp. 6-13, doi: [10.1145/572230.572231](https://doi.org/10.1145/572230.572231).
- Bynum, T.W. (2006), "Flourishing ethics", *Ethics and Information Technology*, Vol. 8 No. 4, pp. 157-173.
- Bynum, T.W. (2018), "Computer and information ethics", in Zalta, E.N. (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 201)*, Metaphysics Research Lab, Stanford University.
- Dama, I. (2017), *DAMA-DMBOK: Data Management Body of Knowledge*, Technics Publications, LLC.
- Ethics4EU (2021), "Research report on European values for ethics in technology", Erasmus+ Project, available at: <http://ethics4eu.eu/european-values-for-ethics-in-technology-research-report/>
- Flanagan, M., Howe, D.C. and Nissenbaum, H. (2008), "Embodying values in technology: theory and practice", *Information Technology and Moral Philosophy*, Vol. 322.
- Floridi, L. (1999), "Information ethics: on the philosophical foundation of computer ethics", *Ethics and Information Technology*, Vol. 1 No. 1, pp. 33-52.
- Floridi, L. (2014), *The Fourth Revolution: How the Infosphere is Reshaping Human Reality*, OUP, Oxford.
- Floridi, L. and Sanders, J.W. (2005), "Internet ethics: the constructionist values of homo poieticus", *The Impact of the Internet on Our Moral Lives*, pp. 195-214.
- Floridi, L. and Taddeo, M. (2016), "What is data ethics?", *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 374 No. 2083, p. 20160360, doi: [10.1098/rsta.2016.0360](https://doi.org/10.1098/rsta.2016.0360).
- Gibbs, A. (1997), "Focus groups", *Social Research Update*, Vol. 19 No. 8, pp. 1-8.
- Gorden, R. (1992), *Basic Interviewing Skills*, F.E.Peacock.
- Gotterbarn, D. (1991), "Computer ethics: responsibility regained", *National Forum*, Vol. 71 No. 3, p. 26, available at: <http://search.proquest.com/openview/fdd917c9e0dbb6018e73d2e11d53229f/1?pq-origsite=gscholar&cbl=1820941>
- Gotterbarn, D., Wolf, M.J., Flick, C. and Miller, K. (2018), "THINKING PROFESSIONALLY the continual evolution of interest in computing ethics", *ACM Inroads*, Vol. 9 No. 2, pp. 10-12, doi: [10.1145/3204466](https://doi.org/10.1145/3204466).
- Helbing, D., Frey, B.S., Gigerenzer, G., Hafen, E., Hagner, M., Hofstetter, Y., Van Den Hoven, J., Zicari, R. V. and Zwitter, A. (2019), "Will democracy survive big data and artificial intelligence?", *Towards Digital Enlightenment*, Springer, pp. 73-98.
- Hilty, L.M. (2015), "Ethical issues in ubiquitous computing – three technology assessment studies revisited", *Ubiquitous Computing in the Workplace*, Springer, pp. 45-60.
- Jacobs, A.R. and Abowd, G.D. (2003), "A framework for comparing perspectives on privacy and pervasive technologies", *IEEE Pervasive Computing*, Vol. 2 No. 4, pp. 78-84.

- Johnson, D. (1985), *Computer Ethics*, Englewood Cliffs, NJ.
- Kranzberg, M. (2019), *Ethics in an Age of Pervasive Technology*, Routledge.
- Kumar, A., Braud, T., Tarkoma, S. and Hui, P. (2020), "Trustworthy AI in the age of pervasive computing and big data", *2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pp. 1-6.
- Loi, M., Heitz, C., Ferrario, A., Schmid, A. and Christen, M. (2019), "Towards an ethical code for data-based business", *2019 6th Swiss Conference on Data Science (SDS)*, pp. 6-12.
- Macnish, K. (2017), *The Ethics of Surveillance: An Introduction*, Routledge.
- Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S. and Floridi, L. (2016), "The ethics of algorithms: Mapping the debate", *Big Data and Society*, Vol. 3 No. 2, 2053951716679679-2053951716679679, available at: file:///Users/454324/Library/ApplicationSupport/MendeleyDesktop/Downloaded/Mittelstadt_et_al-Unknown-The_ethics_of_algorithms_Mapping_the_debate.pdf
- Moor, J.H. (1985), "What is computer ethics?", *Metaphilosophy*, Vol. 16 No. 4, pp. 266-275.
- Moore, J. (2020), "Towards a more representative politics in the ethics of computer science", *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 414-424.
- Morgan, D.L. (2008a), "Focus groups", *The SAGE Encyclopedia of Qualitative Research Methods*, Sage publications, pp. 352-354.
- Morgan, D.L. (2008b), "Snowball sampling", *The SAGE Encyclopedia of Qualitative Research Methods*, Vol. 2, pp. 815-816.
- Müller, V.C. (2020), "Ethics of artificial intelligence and robotics", in Zalta, E.N. (Ed.), *The Stanford Encyclopedia of Philosophy*, Stanford University, Palo Alto CA, available at: <https://plato.stanford.edu/archives/win2020/entries/ethics-ai/>
- O'Keefe, K. and Brien, D.O. (2018), *Ethical Data and Information Management: concepts, Tools and Methods*, Kogan Page Publishers.
- Rasmussen, L.B., Beardon, C. and Munari, S. (2000), *Computers and Networks in the Age of Globalization: IFIP TC9 Fifth World Conference on Human Choice and Computers August 25-28, 1998*, Vol. 57, Springer Science and Business Media, Geneva, Switzerland.
- Saltz, J.S. and Dewar, N. (2019), "Data science ethical considerations: a systematic literature review and proposed project framework", *Ethics and Information Technology*, Vol. 21 No. 3, pp. 197-208.
- Saltz, J., Skirpan, M., Fiesler, C., Gorelick, M., Yeh, T., Heckman, R., Dewar, N. and Beard, N. (2019), "Integrating ethics within machine learning courses", *ACM Transactions on Computing Education (TOCE)*, Vol. 19 No. 4, pp. 1-26.
- Saumure, K. and Given, L.M. (2008), "Convenience sample", *The SAGE Encyclopedia of Qualitative Research Methods*, Vol. 125.
- Shahriari, K. and Shahriari, M. (2017), "IEEE standard review – ethically aligned design: a vision for prioritizing human wellbeing with artificial intelligence and autonomous systems", *2017 IEEE Canada International Humanitarian Technology Conference (IHTC)*, pp. 197-201.
- Stokes, D. and Bergin, R. (2006), "Methodology or 'methodolatry'? An evaluation of focus groups and depth interviews", *Qualitative Market Research: An International Journal*, Vol. 9 No. 1, pp. 26-37, doi: [10.1108/13522750610640530](https://doi.org/10.1108/13522750610640530).
- Taddeo, M. and Floridi, L. (2018), "How AI can be a force for good", *Science*, Vol. 361 No. 6404, pp. 751 LP-752, doi: [10.1126/science.aat5991](https://doi.org/10.1126/science.aat5991).
- Thornley, C.V., Murnane, S., McLoughlin, S., Carcary, M., Doherty, E. and Veling, L. (2018), "The role of ethics in developing professionalism within the global ICT community", *International Journal of Human Capital and Information Technology Professionals (IJHCITP)*, Vol. 9 No. 4, pp. 56-71.
- Vyakarnam, S. (1995), "FOCUS: focus groups: are they viable in ethics research?", *Business Ethics: A European Review*, Vol. 4 No. 1, pp. 24-29.

Webb, H., Patel, M., Rovatsos, M., Davoust, A., Ceppi, S., Koene, A., Dowthwaite, L., Portillo, V., Jirotko, M. and Cano, M. (2019), "It would be pretty immoral to choose a random algorithm", *Journal of Information, Communication and Ethics in Society*, Vol. 17 No. 2.

Weizenbaum, J. (1976), *Computer Power and Human Reason: From Judgment to Calculation*, W. H. Freeman & Co., San Francisco.

Whittaker, M., Crawford, K., Dobbe, R., Fried, G., Kaziunas, E., Mathur, V., West, S.M., Richardson, R., Schultz, J. and Schwartz, O. (2018), *AI Now Report 2018*, AI Now Institute at New York, NY University, New York, NY.

Zachman, J.A. (2008), *The Zachman Framework: The Official Concise Definition*, Zachman International.

Zikmund, W.G. (1997), *Exploring Marketing Research*, Dryden Press. Fort Worth, TX.

Further reading

Davies, N. (2013), "Ethics in pervasive computing research", *IEEE Pervasive Computing*, Vol. 12 No. 3, pp. 2-4.

Corresponding author

Ioannis Stavrakakis can be contacted at: ioannis.stavrakakis@tudublin.ie