The Ethics of Computing: A Survey of the Computing-Oriented Literature

BERND CARSTEN STAHL and JOB TIMMERMANS, De Montfort University
BRENT DANIEL MITTELSTADT, University of Oxford

Computing technologies and artifacts are increasingly integrated into most aspects of our professional, social, and private lives. One consequence of this growing ubiquity of computing is that it can have significant ethical implications that computing professionals need to be aware of. The relationship between ethics and computing has long been discussed. However, this is the first comprehensive survey of the mainstream academic literature of the topic. Based on a detailed qualitative analysis of the literature, the article discusses ethical issues, technologies that they are related to, and ethical theories, as well as the methodologies that the literature employs, its academic contribution, and resulting recommendations. The article discusses general trends and argues that the time has come for a transition to responsible research and innovation to ensure that ethical reflection of computing has practical and manifest consequences.

Categories and Subject Descriptors: K.4.1 [Public Policy Issues]: Ethics; K.7.4 [Professional Ethics]: Professional Ethics

General Terms: Human Factors, Legal Aspects, Management

Additional Key Words and Phrases: Computer ethics, information ethics, responsible research and innovation

ACM Reference Format:

Bernd Carsten Stahl, Job Timmermans, and Brent Daniel Mittelstadt. 2016. The ethics of computing: A survey of the computing-oriented literature. *ACM Comput. Surv.* 48, 4, Article 55 (February 2016), 38 pages. DOI: <http://dx.doi.org/10.1145/2871196>

1. INTRODUCTION

The increasing ubiquity of computing technologies and artefacts leads to the growing relevance of ethical aspects of computing. Despite decades of research in this field, there is no common understanding of key components of the ethics of computing as perceived and put in practice by the communities of technical scholars and practitioners. This article provides the first systematic and comprehensive review of the literature on the ethics of computing.

A sound understanding of ethics is a key component of the professional status that professional bodies such as the Association for Computing Machinery (ACM) aspire to. Professional computer scientists and related technical experts (e.g., software engineers, data scientists) are interested in the social and ethical consequences of their work to ensure that computing can realize its potential benefits. These professionals also need to understand the forces that shape the social and regulatory environment in which they act. Ethics is a key component that can determine acceptance of

Authors' addresses: B. C. Stahl and J. Timmermans, Centre for Computing and Social Responsibility, De Montfort University, The Gateway, Leicester, LE19BH, UK; emails: {bstahl, jtimmermans}@dmu.ac.uk; B. D. Mittelstadt, Oxford Internet Institute, University of Oxford, 1 St. Giles, OX1, 3JS; email: brent.mittelstadt@oii.ox.ac.uk.

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

2016 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 0360-0300/2016/02-ART55 \$15.00

DOI: <http://dx.doi.org/10.1145/2871196>

new technologies as well as legislative and other responses to new technologies. Furthermore, an awareness of social and ethical consequences is increasingly required as a component of funding proposals in which impact statements ask computer scientists to engage with possible and likely consequences of their work. In order to help computer scientists to better understand these issues, we use this article to identify key ethical issues, relevant technologies, applicable ethical theories, and other aspects that characterize current understanding of the ethics of computing. On this basis, we identify gaps in knowledge and link to parallel discussions in other fields. This allows us to recommend research agendas and practical steps to identify and address ethical questions early in research and innovation in computing.

A key stakeholder group that has a crucial influence on the way in which computing research is undertaken and, as a consequence, how computing artefacts and products are developed is represented by computer scientists and related technical experts (e.g., software engineers, data scientists). These experts often determine the technical choices that have a strong influence on ethical and social consequences of computing. At the same time, these technical experts are usually not experts in the normative aspects of computing and in how to assess and evaluate the consequences of their work. This has long been recognized; many attempts have been made to raise the awareness and interest of computing experts in social and ethical aspects of their work, for example, by including it in standard curricula or professional accreditation.

As a consequence, most computing experts who have gone through a structured training program such as a university degree course have an understanding of professional commitments to ethics as represented in codes and expectations of professional bodies such as the ACM, the British Computing Society (BCS), the Institution of Engineering and Technology (IET), and others. This focus on professionalism is important to help embed ethics into professional practice, but it says little about which technologies are of particular interest, which ethical issues are most relevant, or how such issues could be identified or addressed in practice.

The potential of computing technologies to raise ethical and social issues that differ fundamentally from those raised by other technologies has been discussed since the very inception of digital computing [Wiener 1954]. This has led to a steadily growing academic discourse engaging scholars from numerous fields and disciplines. We are now at a point at which computing technologies and devices pervade most aspects of personal, organizational, and social life. As a consequence, the ethical and social consequences of such technologies are being examined with increasing vigor. Scholars with an interest in the ethics of computing come from a broad range of fields including computing, philosophy, law, and social sciences. This article surveys the literature on the ethical and social consequences of computing that is available in mainstream academic publications to provide a comprehensive overview of the state of the debate. It explores the discourse to find out which technologies are investigated, which ethical issues are linked to these, how claims concerning the ethics of computing are supported by research, and which conclusions and recommendations are made on this basis.

While there is much work on the ethics of computing, including several overviews of the field [Brey and Soraker 2009; Bynum 2008b] and anthologies aiming to cover the main topics [Bynum and Rogerson 2003; Johnson 1985; Johnson and Nissenbaum 1995; van den Hoven and Weckert 2008; Himma and Tavani 2008], this article is the first that explicitly explores the discourse as a whole. By focusing on mainstream academic journals, we can understand how the topic of ethics and computing is represented to technical audiences.

In addition to making an important contribution to knowledge by showing what is currently known about the complex relationship between ethics and computing, the

article can provide the basis for policy in both research and education in computing. It clearly highlights which topics are currently most widely discussed and traces the trajectory of topics, issues, and methods employed, suggesting areas requiring further work.

An article offering a survey of a large topic area such as ethics and computing cannot, by necessity, go into much depth with regard to all aspects of discourse that are surveyed. We elaborate on some of the substantive ethical issues (e.g., privacy, professionalism, autonomy, and so on) in the Discussion section. However, the main purpose of the article is to map the topics and issues that have been discussed in the field during the last decade. This is of importance to computing professionals who need to gain insights into how ethical debates relate to their work. It is furthermore of interest to scholars who focus on ethics and computing, and who want to contextualize their work in a broader context.

Section 2 presents an overview of the conceptual background of the ethics of computing. Section 3 presents the survey methodology, including how papers were collected and analyzed. On this basis, we present the findings in Section 4. Section 5 highlights important trends and topics in the discourse. These form the basis of the conclusion in Section 6, which outlines theoretical and policy implications and links the article to the broader emerging context of responsible research and innovation.

2. THE ETHICS OF COMPUTING

This survey provides an overview of the breadth of ethics and computing in the literature. One key audience is that of computing professionals with limited background in ethics. We therefore need to start with an overview of the term. “Ethics” has many interrelated but distinct meanings (see Stahl [2012] for more detail). At the most basic level, it refers to the perception of something being good or right. One may speak of an “ethical use of computing” and mean that it is right, proper, acceptable, or socially appropriate. Such an intuition of the ethical quality of an act is usually based on more or less explicit norms and values that are accepted within a social group or culture. Where such values and norms cease to be easily applicable or where they clash, explicit reflection on the bases and assumptions related to ethical judgments is required. This is what ethics as the discipline of moral philosophy does.

Not all computing professionals have a deep intrinsic interest in understanding the details of ethics. It is therefore important to point out that many of their practical discussions and decisions are nevertheless driven by ethical ideas and principles; thus, it is important to have a general grasp of them. This may be most obvious in cases of moral dilemmas, in which a technical choice is made in the face of competing values. One example could be work on Brain Computer Interfaces (BCIs) that can be used to help severely disabled individuals to better communicate with their environment. While this is doubtlessly a positive moral good, technical choices can affect the ability of users to act in a certain way and thereby impede their autonomy. It is not immediately obvious how particular technical choices in this context could be viewed from an ethical perspective. A different and topical example is that of ethical questions in big data. Large datasets can be used for an enormous range of applications that promise numerous benefits, ranging from improved public health to optimized consumption experiences. At the same time, they raise deep questions about concepts such as privacy and ownership. Identifying the ethical issues in this case is a complex endeavor in its own right; linking them to technical choices that strike an appropriate balance between contradictory interests and values is similarly difficult. The implicit reasoning in such cases can be highly complex. When made explicit, such dilemmas may require recourse to ethical concepts such as utility, virtue, or responsibility. In order to

appreciate the depth of such problems, a certain level of understanding of philosophical ethics is required.¹

Philosophical ethics can be divided into subdisciplines comprising metaethics, normative ethics, and applied ethics [Marturano 2002]. In addition, one can find reference to descriptive ethics as part of moral philosophy [Nijssingh and Duwell 2009]. Descriptive ethics aims to describe and understand moral values, judgments, and practices. This can be the basis of normative ethics, which goes beyond description and seeks to justify particular ethical positions. Normative ethics has developed a number of well-known positions that pervade ethical discourse and that are widely used in reflecting on why a particular action might be considered good or bad. Prominent normative ethical positions or theories include consequentialism, deontology, and virtue ethics. Consequentialism states that the ethical value of an action is to be found in its consequences. The most widely used consequentialist theory is that of utilitarianism, which holds that the aggregate amount of happiness minus the aggregate amount of pain caused by an action are the measure of its ethical quality. These ideas are strongly linked to British 19th-century thinkers such as Mill [1861] and Bentham [1789]. Deontology, on the other hand, holds that the moral quality of an action is to be located in the intention of the agent. Deriving from the Greek work for “duty,” the most famous proponent of deontology, the German philosopher Immanuel Kant [1788, 1797] formulated the Categorical Imperative, which postulates that the hallmark of ethical quality is the universalizability of a maxim. Virtue ethics, the final example to be outlined here, disagrees with the focus on individual action that both consequentialism and deontology share. From the virtue ethics perspective that goes back to Greek Antiquity [Aristotle 2007] but still has strong proponents [MacIntyre 2007], the place of ethics is the character of the agent. Something is ethically good that reflects the temperate character and is expressed in reflective practice.

This very brief introduction to normative ethics cannot do justice to the richness of the positions and centuries of discussion that they have engendered. It suffices for this article, however, in that it gives an insight into some of the most broadly discussed positions that are used to reflect on and justify ethical positions related to computing. It is important to point out that ethical theorizing did not stop in the 19th century and that there are numerous recent normative ethical theories in addition to the three just outlined, such as ethics of care [Gilligan 1990; Adam 2001] or postmodern ethical positions [Sarker et al. 2009].

Ethical theorizing raises numerous fundamental questions, for example, concerning the truth value of moral claims, the question of whether morality is culturally relevant, whether there are moral facts and how they could be known, and whether and how ethics is linked to empirical facts. These questions are the subject of metaethics, which “explores as well the connection between values, reasons for action, and human motivation, asking how it is that moral standards might provide us with reasons to do or refrain from doing as it demands, and it addresses many of the issues commonly bound up with the nature of freedom and its significance (or not) for moral responsibility” [Sayre-McCord 2014].

The final aspect of moral philosophy and the one most relevant to this article is that of applied ethics. Applied ethics uses the ideas and discourses from moral philosophy to come to a better understanding and often to prescriptive positions in particular areas that call for ethical attention. Prominent examples are biomedical ethics, technology ethics, military ethics, or environmental ethics. Applied ethics can include subsections

¹Readers may be interested in more detailed practical examples. This article does not offer the space for such a discussion but can point the reader to repositories of examples, such as the Observatory for Responsible Research and Innovation in ICT: <http://responsible-innovation.org.uk/page-resource/Ethical%20issues>.

of the other types of ethics; thus, there may be descriptive, normative, and metaethics in all of them. For the purposes of this article, the ethics of computing can be seen as a component of applied ethics

A further aspect of applied ethics that is relevant to computing is that of professionalism [Gotterbarn et al. 1999]. Questions regarding the definition and implementation of professions have connotations affecting ethics. One can find discussions of ethical aspects of professionalism in all professional fields (notably medicine and law, but also elsewhere, e.g., in journalism or media) [Gorman 2001; Mostaghimi and Crotty 2011]. Questions regarding the professional nature of the field of computing have long been debated and continue to be a focus for the ethics of computing [Albrecht et al. 2012; Bynum and Rogerson 2003; Estell and Christensen 2011; Martinsons et al. 2009; Payne and Landry 2006; Weckert and Lucas 2013].

2.1. Computer and Information Ethics as Applied Ethics

The term “computer ethics” can be traced back to the 1970s [Bynum 2008b] but ethical concerns arising from digital computing go back to the very beginning of the development of the technology [Wiener 1954; Wiener 1964; Bynum 2008a]. Norbert Wiener, the father of cybernetics and one of the pivotal figures in the development of digital computers, recognized early that this technology had the potential to change many aspects of life that had crucial ethical importance. He foresaw morally relevant changes to the way that we conduct ourselves in many ways, including warfare. His foresight was such that, at the very beginning of the development of computing technology, he explored how computers might change the world of work. Another pivotal moment in recognizing the ethics of computing occurred when Weizenbaum [1977], while working on artificial intelligence applications, found that many people were happy to communicate with computers, some going so far as to accept them as partners in psychological consultations. For Weizenbaum, this was the basis of an exploration of the relationship between computers and humans.

While there are thus early examples of high-level attention to the relationship between computers and ethics, a broader discourse only started in the 1980s and 1990s. During this period, computer ethics developed into a field of applied ethics. Dedicated courses on computer ethics were included in curricula, textbooks on the topics were written [Johnson 1985], and academic conferences (e.g., Computer Ethics Philosophical Enquiry, Computers and Philosophy, ETHICOMP) and journals (e.g., *Ethics and Information Technology*) were created. Professional organizations focusing on computers, such as the ACM or BCS, established ethical codes to provide guidance to their members.

At the same time that computer ethics became recognized as a legitimate field of academic inquiry, the subject area—computers—became less recognizable. The increasing integration of computing artifacts into other technologies and the environment as expressed in visions of technology captured by concepts such as ambient intelligence [Sadri 2011; Friedewald et al. 2005] or ubiquitous computing rendered the idea of computers as easily identifiable artifacts obsolete. This diffusion of computers into their surrounding environment raises new ethical questions related to issues such as privacy, surveillance, autonomy, or ownership [Wright et al. 2008; Quilici-Gonzalez et al. 2010; Lally et al. 2012].

To complicate matters further, there are terms—in addition to computer ethics and information ethics—that cover similar, if not identical, territory. These include ICT ethics [Markus and Mentzer 2014], IT ethics [Banerjee et al. 1998; Miller 2009], cyberethics [Spinello and Tavani 2001], or ethics of information systems [Cohen and Cornwell 1989; Mingers and Walsham 2010], to name some of the more prominent ones. We therefore need to delineate the subject area of this article.

Fig. 1. Possible components of an article covering ethics of computing.

2.2. The Ethics of Computing

This article surveys the current state of the academic discourse on ethics and computing. The outline of computer ethics is important to understand a possible conceptual basis for it, but it is by no means certain that the technology-oriented communities working on computers are aware of, and engage with, computer and information ethics discourses.

At the same time, an awareness and understanding of ethical issues of computing is gaining importance. Computer ethics may have been a very specialized area of interest during its inception several decades ago. The ubiquity and pervasiveness of computing devices have changed this radically. Current political and social debates about (state) surveillance, “Big Data,” intellectual property of digital content, social consequences of widespread use of social media, and many more indicate that ethical issues related to computing have taken center stage.

The approach that we take is to explore the literature on the ethics of computing that is available to technical research communities to understand the subject area from their side. This was achieved by exploring the discussion of ethics and computing in mainstream academic outlets as represented in the databases Scopus and ISI Web of Knowledge, to which academic users normally have access. Instead of defining *a priori* what counts as a computer, an ethical issue, or an ethical theory, we have read the literature on ethics and computing with a view to determining the meaning of these concepts in published research. We use the term “ethics of computing” to underline our initial openness with regard to these issues.

There is a diverse body of literature that covers many angles of ethics and computing. The survey of literature described in this article aims to establish the important components of this body of literature. One starting assumption underpinning this article

is that published research on ethics and computing can have the following structure: (1) a paper will normally focus on one or more ethical issues (2) that are often linked to particular technologies and (3) can be assessed from the perspective of particular ethical theories. Such research may (4) use identifiable methodologies to come to insights. These are likely to (5) constitute specific contributions to the literature and (6) can warrant a range of recommendations.

We realize that this structure of a paper on the ethics of computing is unlikely to be reflected fully in each piece of research. However, we believe that it is a useful way of understanding how ethics of computing can be conceptualized. The key component here is the “ethical issue”—or the thing, relationship, or process—that is seen as ethically problematic. The five other components establish the context in which the ethical issue is given relevance and meaning. This structure allows an understanding of a paper in the field. Moreover, it provides our research with a structure that allowed for the development of the methodology described in the following section.

3. METHODOLOGY

Understanding the literature is a crucial aspect of all research. Papers reviewing the literature are key to consolidating existing knowledge, identifying gaps in current knowledge, and developing research agendas. Literature reviews create structure within a discipline by identifying a collective representation of what is known and what needs to be known [Boell and Cecez-Kecmanovic 2014]. According to Rowe [2014, p. 242], one can distinguish between different goals of literature reviews:

- (1) to summarize prior research
- (2) to critically examine contributions of past research
- (3) to explain the results of prior research found within research streams
- (4) to clarify alternative views of past research (not necessarily integrated together)

According to the respective aim of the review, there are different approaches and styles. Two key styles are statistical meta-analyses, which quantitatively assess the discourse and prior reviews as a whole, and narrative reviews, which rely on qualitative analysis to gain a deeper understanding of particular issues. Some studies combine the quantitative and qualitative analysis to provide deeper insights [Joseph et al. 2007]. Such an approach is taken in this article, which undertakes a systematic review that differs “... from traditional narrative reviews by adopting a replicable, scientific and transparent process, in other words a detailed technology, that aims to minimize bias through exhaustive literature searches of published and unpublished studies and by providing an audit trail of the reviewers [sic] decisions, procedures and conclusions” [Tranfield et al. 2003, p. 209] This section describes the methodology employed in selecting appropriate literature and identifying and coding themes, as well as the processes used to validate the quality of the review.

3.1. Selection of the Literature

As the focus of this article is on the ethics of computing as perceived by computer scientists and members of related communities, the starting point of our literature search was to focus on outlets that would be available to such communities. It was therefore decided to focus on widely available academic databases that technical scholars are likely to have access to. In order to give some breadth to the search while simultaneously supporting a focused approach, we searched popular databases Scopus and ISI Web of Knowledge. These databases cover the majority of technical publications and index other computing-specific databases such as those of the ACM or IEEE. Both Scopus and ISI include subject or research areas; we chose to focus on computing in our search. Google Scholar was excluded due to the exceeding breadth of its

coverage and the difficulty of focusing on computing-specific publications. Using two databases ensured that possible gaps in coverage in one could be compensated by the other.

Given the rapid pace of technological development in computing, we restricted our focus to recent developments that represent issues of contemporary importance. We therefore included one decade of publications, from 2003–2012. While conferences have the status of important and legitimate outlets in computer science and related disciplines, we nevertheless limited our sources to journal publications under the assumption that strong conference papers are likely to be turned into journal papers. Additionally, journals typically have more stringent peer review and quality control.

Having thus decided on the range of publication dates (2002–2012) and type (journal) to be included, we searched the Scopus and ISI databases using the search term “ethic*.” The wildcard in the search term ensured that a range of topics and titles would be included. The search was run on January 17, 2013, leading to 1597 hits on Scopus and 1296 hits on ISI.

The responses from both databases were collected, duplicates were removed, and papers to be included in the full analysis were identified. In order to ensure the comprehensiveness of the analysis, each author suggested 10 papers that they saw as important in the field. The sample was then checked to ensure that these papers were represented. This initial validation of the sample highlighted some minor issues with the databases, for example, the fact that certain journal issues were labeled as conference papers. These mistakes were manually rectified.

Using the validated set of literature, we then looked at each paper to ensure its relevance to the survey. Criteria for inclusion were:

- relevance to the ethics of computing of the paper as a whole;
- ethics of computing applications constituted the major topic of the paper;
- coverage of well-discussed ethics issues, for example, surveillance, privacy, or professionalism in computing;
- clearly ethical topics with plausible relevance to computing;
- teaching computer ethics;
- Internet research ethics.

Papers were expected to fulfill one or more of these criteria. Papers not specifically on the ethics of computing were removed from the sample. Exclusion criteria were:

- basis in fundamentally different discipline (e.g., pharmacogenetics);
- predominant focus on business ethics;
- ethical issues for which information/computing are incidental (e.g., communication with cancer patients);
- religious issues;
- ethics mentioned only as a reference, but not clearly linked to the paper;
- teaching computing or computer sciences;
- research ethics that refers to technology only in general.

The title and abstract of each paper was analyzed by one of the authors to assess whether it clearly fell into the specified criteria. In cases in which this was not obvious, all three authors looked at the paper and, when necessary, read it to assess relevance. At the end of the selection process, 809 papers remained in the sample. During the coding process (details to follow) and in light of the more detailed understanding of the papers that we developed, further papers were removed from the sample for not satisfying the inclusion criteria. The final number of papers coded was 599. Bibliographic information for the whole sample is available in the supplementary material.

Fig. 2. Temporal distribution of papers (n = 599).

Figure 2 shows the temporal distribution of the papers. An increasing trend is clearly visible, underlining the growing importance of the ethics of computing.

3.2. Identification and Coding of Themes

The aim of identifying the discourse on ethics and computing and the intricacies of ethical arguments called for a mixed-methods approach [Mingers 2001] that included a qualitative analysis of the content of the paper in conjunction with a quantitative analysis of these qualitative insights. To this end, we uploaded all full papers in PDF form to NVivo Server, version 10. NVivo (http://www.qsrinternational.com/products_nvivo.aspx) is a tool typically used for the analysis of qualitative data while remaining supportive of work with quantitative data. In order to understand the discourse on ethics and computing, we coded each paper using a scheme based on the model in Figure 1. For each paper, we aimed to identify at least one (1) issue (i.e., ethical question or problem at the core of the paper), (2) technology, (3) ethical theory, (4) methodology, (5) contribution and (6) recommendation.

Coding was done by reading the papers and highlighting relevant words or sentences in the paper and coding them under the top-level code listed earlier (i.e., issue, technology). Wherever possible, the original wording was used as the code. In cases in which it was necessary to paraphrase the code, we tagged this explicitly. While almost all papers discussed an explicit issue(s), the presence of other items was often less clear. Multiple codes in the same paper from a single type were possible. For example, a number of papers discussed several issues or used more than one theory. When a code for a particular item could not be identified, it was coded as “none.” Coding stayed as close to the text of the papers as possible, with only items coded that were recognizably intended by the paper’s author(s). For example, if a paper did not have an explicit statement about its contribution, we refrained from giving our interpretation of what the contribution might be and instead coded it as “none.”

The quality-assurance process for the coding was similar to that for selecting relevant literature. Initially, all three authors separately coded the same sample of 30 papers and discussed their coding approach. Having agreed on consistent principles, the remaining papers were then divided between the three authors and coded individually. When the coder came across problems or questions, the papers were set aside and all three authors discussed each to decide upon appropriate coding. Table I shows how many unique codes were created for each of the six coding types.

The numbers of codes reflects the diversity and frequency with which codes were created within each type. Concerning frequency, of all the types of issues were coded most frequently, with many sources having multiple issues coded. Differences in diversity reflect the characteristics of the type of code. Methodology, for instance, has a rather

Table I. Unique Codes Per Type

Item	Number of Unique Codes
Issue	1133
Technology	354
Ethical Theory	308
Methodology	96
Contribution	229
Recommendation	500

bounded set of possible instantiations, such as surveys, experiments, or case-studies, while technology has a broader set of options.

3.3. Categorization of Codes

The final step in coding was the categorization of codes. In the case of issues, for example, our sample had a total of 1133 codes. Many of these were similar and overlapping due to naming the codes *in vivo*, using the words of the paper's author(s). For the purpose of this review, it is not helpful to have a long list of similar codes; instead, these needed to be grouped into useful categories to allow for a more detailed understanding of past and current questions in ethics of computing. In order to structure the findings, we aimed to have a relatively small number of groups under each code type (i.e., Technology, Issue). To do so, the codes assigned to the 30 papers cross-coded for validation (provided earlier) were categorized by all three authors. As with the initial coding process, categories were compared and aligned through discussion to ensure consistency. Having established the main categories, all of the codes were split between the authors, who each categorized a segment individually. As before, contentious or questionable categorizations were highlighted and agreed upon by all authors.

Categorization produced a hierarchical database of codes of up to four levels (see: Figure 3). The hierarchy is structured as follows: for each type of code (e.g., Technology, Issue, Contribution), the specificity of codes and categories increases at each subsequent level. The bottom level of the hierarchy (Level 3 or 4, depending on the type of code) contains the raw codes from the reviewed papers. Each level consists of categories describing the codes.

Having categorized all the codes, the final step of the review was to analyze our codes and categories to describe key findings, trends, and developments. These are described in the next section.

4. FINDINGS

As outlined earlier, the logic of our analysis was that a paper on the ethics of computing could cover one or several ethical issues, most likely linked to one or more technologies. These issues could be evaluated from the perspective of an ethical theory. We furthermore expected papers to describe the methodology used to arrive at their insights, to potentially spell out their contribution to knowledge, and finally to make explicit the recommendations arising from them.

When coding the papers for their contributions to these categories, it emerged that these contributions were unevenly distributed. Almost all papers discussed issues, and most covered technologies. A large number made explicit use of ethical theories, but methodology, contribution, and recommendation were not components of all papers. Whenever a novel item in any of these categories was found (e.g., an ethical issue not yet covered or a new technology), it was coded as a new code.

Figure 4 represents a word cloud of all papers that were retained for full analysis ($n = 599$). It provides interesting indications of some of the main technologies, issues and aspects of ethics in computing that will now be discussed in more detail.

Fig. 3. Example of hierarchical categories for technology codes.

In this section, we describe the findings of the survey broken down with regard to the main aspects of this type of argument. Table II provides an overview of the 20 items (excluding “None” and “Other” codes) found most frequently in each of the four code types of Technology, Issue, Ethical Theory, and Methodology². For each of these we show the name of the item and the number of papers coded.

Findings for each code type are presented throughout this section at the level representing the best balance between specificity of codes and meaningful groupings (Level 2 or 3). To aid the reader in identifying the level at which the discussion is operating and the relationship between codes within a type, interactive mind maps have been provided for each code type (see: Appendix 2). The ideal level of discussion varies for each code type according to the total number of codes and the relative simplicity of the raw codes. For example, raw Methodology codes are meaningful at Level 3 due to substantial standardization in methods terminology within the reviewed literature, whereas the specificity of Contributions means that meaningful discussion is possible only at Level 2 categories.

We start the following discussion with ethical issues, as we believe these to be of most interest to the readers of this article. We provide a brief discussion of some of the key issues to allow for a better understanding of what they mean in a computing context. The next section looks at the technologies discussed in the discourse. From there, we move to ethical theories, methodologies, contributions, and recommendations

²The Contribution and Recommendation types did not provide sufficient codes to construct a top 20, as seen in Sections 4.5 and 4.6.

Table II. Top 20 Items for Each Code Type (Levels 2 + 3)

Technology		Issues		Ethical Theory		Methodology	
Internet	104	Privacy	177	Information	35	survey	40
ICT	84	professionalism, work-related issues	70	consequentialism	30	literature review	22
software	54	autonomy	58	ethical codes	29	questionnaire	16
robots	46	Agency	49	Deontology	25	interviews	15
computing	34	Trust	49	Kant	21	quantitative	14
artificial intelligence	32	specific computing technologies	47	Justice	18	case study	13
e-Health	25	Consent	43	Floridi	15	scenario	8
e-services	22	Identity	42	human rights	13	case studies	6
social networking	21	inclusion, digital divides	42	Aristotle	12	content analysis	5
artificial agents	20	Security	41	responsibility	10	grounded theory	4
information systems	20	harm, misuse, deception	39	design, innovation, technology	10	qualitative	4
bio-ICT	18	health-related issues	38	Asimov	10	focus groups	4
computer surveillance	13	justice	32	legal	9	ethnography	4
ambient intelligence	12	behavior, practical issues	31	autonomy, freedom	9	experiments	4
electronic health records	11	epistemological issues	28	robot, AI	9	review	3
computer games	11	design	26	principalism	9	book review	3
peer-to-peer	10	education-related issues	24	Well-being, human good	8	conceptual	3
Web 2.0	9	research ethics	21	medical ethics	8	laboratory experiment	3
Virtual reality	9	moral values and duties	21	Rawls	8	modelling	3
e-commerce	9	online piracy	20	critical	8	case examples	2

4.1.1. Core Ethical Concepts. We used the category of “core ethical concepts” to collect all those issues that reflect a significant aspect of the broader discourse in moral philosophy. The dominant ethical issue that by far outweighed any other ethical issue in this or any other category was that of privacy. A total of 180 papers from our sample discussed privacy as the, or at least one of the, central ethical issues of computing. The second most widely discussed issue was autonomy ($n = 59$), followed by agency and trust (both $n = 49$). Figure 6 shows the distribution of the core ethical issues in our survey by number of papers coded for each category, followed by relative distribution.

These core ethical issues are characterized by the fact that they constitute important topics of ethical reflection. At the same time, they are not necessarily linked to particular ethical theories or position. Freedom and justice, for example, have been discussed and applied in philosophy and ethics extensively, with numerous competing theoretical positions seeking to define the meaning of these terms. For this reason, we distinguished these core ethical issues from issues that pertain directly to ethical theory or other theoretical philosophical questions.

Fig. 5. Main categories of ethical issues, development over time (Level 2).

Fig. 6. Core ethical issues distribution (Level 3).

4.1.2. Fundamental and Theoretical Issues. Fundamental and theoretical issues are generally those that require a better conceptual and theoretical understanding in order for ethical questions to become clear or to be amenable to solutions. The largest subcategory in this area is that of epistemological issues ($n = 30$), that is, related to the question of what is knowledge. This finding reveals that ethical judgments need to be based on an understanding of the workings of the social and natural world, and how we form reliable knowledge about it. Lack of clarity in such matters of fact can raise problems with regard to ethics. Bias in research on computing, for example, can skew not only the factual understanding but also the way in which ethics is perceived.

Another set of issues related to the question of epistemology and ethics is that of ontology. We found 13 papers that highlighted ontology as an ethical issue. We use the term “ontology” in the philosophical meaning, that is, as referring to the nature of being. There is a complex relationship between what is, what can be known or perceived, and what is seen as morally good or justifiable. While this is a well-established fact in moral philosophy, the majority of papers referring to ontology did so in the context of one ethical theory: Floridi’s information ethics [Floridi 2010].

The second largest group of theoretical issues referred to particular moral values and duties ($n = 21$). Examples of duties discussed in the sample are friendship [McFall 2012], honesty [Mostaghimi and Crotty 2011], or sustainability [Williamson 2010]. Further noteworthy groups of fundamental and theoretical issues include machine

Table III. Social and Practical Issues

Issue	#
age-related issues	4
behavior, practical issues	31
codes and policies	16
consent	43
cultural issues	20
economic and business issues	19
education-related issues	24
environment, sustainability	2
harm, misuse, deception	39
health-related issues	47
inclusion, digital divides	45
military	13
other	29
politics	20
power	12
professionalism, work-related issues	70
quality of life—welfare	14
research ethics	23
society, impact on	9

ethics, issues arising from specific ethical theories, moral status of technology and machine ethics, or moral principles. What all of these issues have in common is that they are related to complex philosophical discourses and not open to simple answers. All relate directly to questions of ethical theory (see Section 4.2). They tend to be contested within philosophy and do not lend themselves to algorithmic solutions or practical applications. Fundamental and theoretical issues thus stand in stark contrast to more practical issues of immediate importance to a specific context.

4.1.3. Social, Practical Issues. The largest group of ethical issues consists of those that we classed as social and practical issues (see Table III). In this group, we included all those issues that have implications for the way in which humans live together and interact. This group of issues was discussed in 331 papers, more than half of the literature we investigated.

The topics that fall under this category cover many of the main issues that have long been of interest in ethics and computing. It is worth noting that the largest category here is “professionalism and work-related issues,” which confirms the continuing importance of professionalism in matters of ethics and computing. The relationship between computing and health is another key issue, as is the question of inclusion and digital divides. This category furthermore covers practical questions consent, which can be important both in research and in commercial and professional practice, in which personally identifiable data is involved. Overall, the broad spread of these nodes shows the breadth of influence that computing has on all aspects of life. Computing can affect quality of life, power and politics, economics, education, the environment, and virtually all other aspects of life. This breadth underlines the importance of ethics in computing, but also highlights the difficulty of gaining a clear understanding of the rich array of ethically relevant social consequences of computing, many of which may be entangled with legal and regulatory issues.

4.1.4. Legal Issues. We described as legal issues those pertaining to existing or future formal regulation, in most cases, statutory regulations and law. The relationship between law and ethics is complex. Ethics and shared moral beliefs shape the law to

a significant degree. At the same time, legal regulations can enforce, strengthen, or modify moral positions. However, this relationship is not always straightforward and there are numerous examples in which legislation and morality do not converge or in which the moral evaluation of the law is contested. In the field of computing, this is most notably the case with regard to questions of intellectual property (IP). The debate concerning who does or should own the intellectual property in, for example, software or digital content is one of the cornerstones of discourse in computer ethics [Hull 2012; Spinello and Tavani 2004]. It is therefore not surprising that the majority of issues in the category of legal issues can be subsumed under the heading of intellectual property.

We broke down these codes further to distinguish which aspect of IP that they addressed. We used two broader codes, namely “intellectual property” ($n = 20$) to cover all IP issues and “ownership” ($n = 14$) to cover property issues going beyond IP. Key nodes pointing to more specific legal aspects were “copyright” ($n = 8$) and “online piracy” ($n = 20$).

Another ethical issue with important legal ramifications is the dominant ethical issue of the survey: privacy. Privacy can be seen as a moral right that is sometimes expressed in law and is even upheld as a human right, for example, in Article 8 of the European Convention of Human Rights. One way of implementing this right is via data-protection legislation, for example, through the European Directive 95/46/EC which provides the basis for national data-protection legislation of EU Member States. For the sake of simplicity, however, we did not code privacy and data protection separately, but grouped all papers referring to the ethical or legal side of this issue under “privacy.”

4.1.5. Technical Issues. The final category of codes consisted of all those ethical issues that are caused or linked to particular technical functions or artifacts. The top issue in this category was security ($n = 41$), followed by design ($n = 26$). The former specifically addressed technical means of security, the adequacy of which can determine the ethical acceptability of a technology. The latter addressed how design can affect ethical issues or how it can be shaped by ethical concerns. A key idea here is that of value-sensitive design [Cummings 2006; van Wynsberghe 2013; Friedman et al. 2008], a methodology that incorporates value considerations early in the design process.

In addition to these two technical issues, numerous references were found to ethical issues specific to particular technologies. The two largest of these were robots ($n = 23$) and the Internet ($n = 7$), which are discussed in detail in the following section.

4.2. Technology

In many cases, ethical questions arise in the context of a particular technology. We therefore expected papers to explain to which technology the ethical aspect was particularly relevant. As in all other cases, we used a bottom-up coding logic, that is, we coded the technologies as we found them in the papers. These codes were subsequently categorized by grouping technologies based on shared technological capacities or intended uses. Grouping allowed for a better understanding of which types of technologies were of most interest in the reviewed literature, rather than relying on authors using the same terms to describe closely related technologies (e.g., consider the description of an unmanned aerial vehicle as an “autonomous weapons system” or “autonomous warfare”). Figure 7 shows the broad distribution of technologies at the top level of grouping (Level 2, per Section 3.3).

While many technologies easily fit into a recognizable category, the same did not hold for all. The two largest top-level categories, “generic ICT” and “other technologies,” are rather broad. However, a look at the content of these categories shows that they consist of generic descriptions of ICTs. The category of “generic ICT” included codes such as “computing” ($n = 37$), “ICT” ($n = 86$), “software” ($n = 54$) and “information systems”

Fig. 7. Technology codes distribution (Level 2).

($n = 21$). This indicates that a large number of papers discuss the ethics of computing in rather broad terms without reference to specific technologies.

However, there were a number of clearly identifiable technologies or technical themes in the sample. Internet and related technologies were a main focus. This includes generic references to the Internet ($n = 105$) as well as specific Internet-enabled technologies and applications such as e-commerce ($n = 9$), e-research ($n = 7$), or e-services ($n = 22$). Health-related technologies formed a similar, clearly identifiable cluster consisting of “e-health” ($n = 25$), “electronic health record” ($n = 11$), “bio-ICT” ($n = 29$), and further related technologies. The three most clearly defined technologies were “artificial intelligence” ($n = 68$), “robotics” ($n = 56$) and “social media” ($n = 40$).

These findings suggest that numerous high-level technologies attract considerable ethical attention (AI, health-related, Internet, robotics, and social media). However, these are still very broad categories and are less well represented than the even more general categories referring to ICT as a whole. In order to shed some light on the details of the technologies being discussed, we looked at the technologies discussed in the papers at a more granular level. As in the case of ethical issues explained earlier, we did this by moving from the higher-level categories described earlier and analyzing which technologies were included in them. Due to the large number of technologies included in the survey, our data structure included 4 levels of abstraction. We used this structure to look at the more granular level to see whether it could reveal any further insights into important and widely discussed technologies. Going to Level 3, we found the top 20 technologies listed in Table II. An interesting observation resulting from this visualization of the coding is that technologies tend to be treated quite broadly and generically in discussing ethical issues of computing. Our initial assumption that ethics in computing tends to focus on particular information technologies therefore could be only partially upheld (see Figure 8).

4.3. Ethical Theory

The third component of our hypothesized structure of a paper on ethics and computing is that of ethical theory. Ethical issues may be evaluated differently depending on the theoretical assumptions. We therefore sought to find out which ethical theories were used in our sample. When coding for ethical theory, we aimed to focus on theories that are actually used in the argument of the paper. Brief references to ethical theories without further relevance to the narrative were not coded.

The first interesting observation arising from this is that half of the reviewed papers did not make explicit use of an ethical theory or concept ($n = 300$). The remaining half covered a broad range of theoretical positions (Figure 9).

The largest group, “ethical theories linked to individuals,” covered a number of ethical theories famously linked to specific philosophers, such as Kant ($n = 21$) or Aristotle

Fig. 8. Top 20 technology codes (Level 3).

Fig. 9. Categories of ethical theories distribution (Level 2).

($n = 12$). The most prominent contemporary philosopher in this group is Luciano Floridi ($n = 15$). “Ethical theories linked to individuals” is closely linked to the second-largest category, “traditional ethical theories,” which included ethical theories outlined earlier, for example, consequentialism ($n = 30$), deontology ($n = 25$), and virtue ethics ($n = 7$). Others within this category included principlism ($n = 9$), religious ethics ($n = 3$), and ethics of care ($n = 2$).

In a number of papers, ethical theory was addressed using recourse to generally accepted moral principles, but without specific reference to a particular theory. Salient examples of this include justice ($n = 18$), human rights ($n = 13$), responsibility ($n = 10$), and respect for autonomy ($n = 9$). Unsurprisingly, numerous references were also made to ethical theories that have been developed with the specific view toward computing. These notably include information ethics ($n = 35$) and ethics of robots or AI ($n = 8$). Other theories or sets of principles from applied ethics include ethics of design ($n = 10$), medical ethics ($n = 8$), and research ethics ($n = 6$).

Fig. 10. Top 20 methodology codes (Level 3).

4.4. Methodology

Having thus understood the usage of ethical theory across the sample, the fourth element of the model concerned how claims were supported. We assumed that papers that claim to say something novel on the topic of ethics and computing would ideally substantiate these claims, and that the traditional academic way of achieving this is to establish, justify, and make use of a particular methodology.

Again, our assumption about the structure of a typical paper (see Figure 1) was proven wrong, as 422 of the papers did not mention a methodology used. The distribution of those that did can be seen in Figure 10.

One interesting observation arising from Figure 10 is that the majority of the methodologies mentioned are used in the social sciences. This may reflect the background of the authors and a particular penchant of social sciences for the use of methodology. In our experience, this focus on methodology is much less pronounced in other disciplines, including some of those interested in ethics of computing such as philosophy or computer science. The limited use of the term “methodology” and the dominance of social sciences in the methodologies used may just reflect the academic provenance of the contributing authors.

4.5. Recommendation

If papers on the ethics of computing are meant to make a practical difference, then it is reasonable to assume that they would spell out how this practical difference is to be achieved. To put it differently, one could expect such papers to contain recommendations for the various stakeholders involved.

This assumption proved to be broadly correct, with approximately two-thirds of the sample ($n = 375$) providing recommendations. The distribution of recommendation types can be seen in Figure 11. It is notable that the largest set of recommendations referred to the research system itself, for instance, in endorsing further research on the topic. Practical guidelines, tools, and examples of good practice formed the second largest set of recommendations, followed by awareness raising, contribution to policy and debate, discussion, and engagement.

4.6. Contribution

The final element of the papers that we explored concerned perceived contributions. Any paper published in an academic journal can reasonably be expected to make a

Fig. 11. Recommendation code distribution (Level 2).

Fig. 12. Contribution code distribution (Level 2).

contribution to knowledge. We wanted to identify whether authors had reflected on this contribution and made explicit what they thought it was.

Again, this assumption proved to be misleading, with two-thirds ($n = 409$) of the sample not referring explicitly to their contribution. Where contribution was discussed explicitly, the largest percentage focused on contribution to theory ($n = 57$), followed closely by practical guidance on ethical problems ($n = 54$) and clarification of ethical issues ($n = 53$) (see Figure 12). Given the common assumption that ethical issues arise in connection to specific computing technologies, the relative lack of practical guidance in perceived contributions is somewhat surprising.

5. DISCUSSION

The Findings section shows that the discourse on ethics and computing is very rich. Much discussion exists of issues and technologies and, to a lesser degree, of ethical theory. Methodology, recommendations, and contribution are less widely explicitly discussed. In this section, we now explore in more detail the relationship between these different categories. We begin with an overview of the most dominant relationships between key categories. This is followed by an introduction to nine of the ethical issues featured most frequently in the discourse. For each, we briefly discuss their content, then outline links to other categories and issues. This aims to provide the reader with

Fig. 13. Mapping of connections between top 20 codes for each type.

a better understanding of why the issues are important. The section finishes with a high-level review of the findings that highlights key findings, including gaps in current research.

Section 4 highlighted the top 20 codes for each of the six code types (e.g., Issue, Technology, Methodology). Further insight into the nature of the discourse can be gained by mapping the connections between these high-level codes (Figure 13). For this purpose, we first established the relative strength of the connection between two codes, based on the number of papers in which the codes co-occurred. In Figure 13, a higher strength of a connection is represented by nodes with a greater width and opacity. For example, the mapping shows a strong connection between privacy and the Internet, and e-health and health-related issues.

Second, implementing the Wakita and Tsurumi algorithm [Wakita and Tsurumi 2007], clusters were established based on the similarity of the pattern of connections between codes within the sample. Codes within a cluster are better connected to each other in terms of being mentioned in the same source paper(s) than codes from other clusters. For example, the codes “Asimov” and “dilemma analysis” share connections to each other and to “privacy.” In our sample, seven clusters were constructed for the vast majority of codes, with only seven codes not fitting into any of the seven clusters.

In Figure 13, codes that reside in the same cluster are represented in the same color and located in the same box. Each cluster is not independently meaningful, instead providing structure to the figure, which highlights codes that appear together most frequently. The following section provides an introduction to some of the issues with the strongest connections in Figure 13.

5.1. Relationship Between Key Ethical Issues and Other Categories

In line with our contention that ethical issues are the core category of the discourse on ethics and computing, we sought to gain a deeper understanding of the content of these issues and their relationship with other categories. We therefore looked at the top issues as listed in Table II, and looked in more depth at ten of these that represent clearly identifiable discussions. In order to give the reader a working knowledge of how these concepts are used within discourses around the ethics of computing, a brief introduction is provided, indicating common usage and meaning within the reviewed literature, and linkage to other categories.

5.1.1. Privacy. Privacy proved to be the dominant ethical concern and concept employed throughout the discourse ($n = 177$). Its frequency points to a diverse theoretical history and applications to technologies across decades, preceding “computer ethics” as a discipline (e.g., Westin [1970]). To start, privacy can be seen as intrinsically morally valuable, or derive its value from related rights or ethical concepts, often autonomy or freedom (e.g., Schoeman [1984] and van den Hoven [2008]). In relation to computing technologies, two forms of privacy can be broadly, albeit crudely, distinguished: data privacy and personal privacy. The former concerns data about oneself, for which acceptable norms of authorship, movement, and modification can be defined. The latter includes forms of privacy not explicitly addressing data about oneself, such as a “right to be left alone,” distinctions between “public” and “private” spaces and contexts (see Nissenbaum [2004]). As revealed by modern debates over public and private digital spaces and modes of communications (e.g., Zimmer [2010]), the distinction inevitably blurs in technologically mediated interactions. Still, the distinction remains helpful in ontological terms, distinguishing between privacy as control over informational representations of the self and related phenomena (see Floridi [2011, 2013]), and as restrictions on access to physical bodies and spaces [Bagüés et al. 2010].

Numerous standout examples of the variety of issues and mechanisms that can be described in terms of privacy were evident in the reviewed literature. Data privacy can be enacted through privacy policies that act as an agreement between user and operator to maintain confidentiality by restricting processing of personal data [Mizani and Baykal 2007]. Particularly strict norms are often found for computing technologies involved in health care or processing of medical data [Mack 2004]. In protecting privacy norms, an unavoidable link with security exists, as technical measures such as anonymization and security architectures can help ensure that unauthorized access and reidentification of personal data are avoided [Cary et al. 2003; Al Ameen et al. 2012]. In defining appropriate norms for particular user-operator relationships, a tension may be evident between the user’s expectations of privacy and the potential commercial value of the data for the operator. Such concerns are particularly relevant for computing technologies that collect data of which the user may not be aware, including websites and radio-frequency identifiers [Wasieleski and Gal-Or 2008], or for surveillance technologies used to enhance security in public spaces [Bowyer 2004].

This brief overview of example applications of privacy in the discourse indicates the seemingly limitless applicability of the concept. As evidence of this, privacy was discussed in the same context as numerous other ethical concepts and issues in the discourse, including autonomy ($n = 26$), consent ($n = 33$), harm ($n = 15$), health-related

issues ($n = 21$), identity ($n = 16$), digital divides ($n = 17$), professionalism ($n = 26$), research ethics ($n = 15$), security ($n = 19$), trust ($n = 18$), theories of information ($n = 9$), human rights ($n = 8$), and consequentialist theory ($n = 7$). Outstanding issues represented by these figures will be highlighted in the following sections. Further suggesting its seeming ubiquity and relevance over time, privacy was discussed in relation to each technological group coded, in particular, for Internet technologies ($n = 41$). Despite its illustrious past, it would appear that computer ethics has far from exhausted the potential uses and implications of privacy for emerging computing technologies such as quantum computing or the Internet of Things.

5.1.2. Professionalism. Professionalism is a key concern for the ethics of computing. As indicated earlier, there is a long stream of activities by computing-related professional bodies in establishing computing as a profession [Weckert and Lucas 2013]. One core question is whether computing can be considered a profession. One typical defining feature of professions is their responsibility towards the public. Attention to ethics is an underlying justification for granting professional status [Mostaghimi and Crotty 2011]. Describing computing as a profession is thus linked to the expectation that computing professionals will pay attention to ethics. Professionalism counts as an ethical issue in this paper insofar as it has material consequences for the way that ethics of computing is dealt with in the real world.

The literature on professionalism discusses a number of ways of implementing and practicing it. This includes pledges [Albrecht et al. 2012], professional bodies, regulations, and many others. A key ingredient of professionalism is a formalized set of expectations towards professionals that are often expressed in the form of codes of ethics or codes of conduct. This relationship between professionalism and codes was confirmed by our research, which had a clear dominance of “ethical codes” as an ethical theory discussed in conjunction with professionalism ($n = 12$).

Other ethical issues that were frequently discussed in papers dealing with professionalism were privacy ($n = 26$), autonomy ($n = 11$), and harm ($n = 10$). This is consistent with the assumption that one role of professionalism is to avoid harm and privacy being the most prominent issue that can lead to harm. Autonomy is important because it is a component of the status of professionals and therefore plausible as a related issue.

As professionalism is a broad topic that covers all areas of computing, it is probably not surprising that the references to technologies in papers dealing with the topic tended to be broad. The three top technologies appearing with professionalism were all broad and generic: ICT ($n = 12$), Internet ($n = 13$), and software ($n = 12$).

5.1.3. Autonomy. Autonomy, or self-governance, “is the ability to construct one’s own goals and values, and to have the freedom to make one’s own decisions and perform actions based on these decisions” [Brey 2005]. It has been defended as an important value for self-realization in Western society, and is strongly related to the concepts of freedom [Salvini et al. 2008], independence [Hildt 2010], and control [Anderson 2005]. This is reflected in the literature, as from a theoretical perspective, autonomy often is discussed in terms of justice ($n = 9$).

When it is discussed in relation to technology, questions arise about possible infringements, as well as possible enlargements of autonomy. On the one hand, technology may lead to paternalism [Kaplan and Litewka 2008], dependence [Stip and Rialle 2005], and loss of control [Allen and Roberts 2010]. At the same time, it may increase independence [Kaplan and Litewka 2008] and enable human enhancement [Allhoff et al. 2010]. Being autonomous is often associated with having control over your own data. With privacy being a prevalent topic in computer-ethics literature, it comes as

no surprise that, of the 58 papers in the sample addressing autonomy, it was most frequently discussed in relation to privacy ($n = 26$).

Autonomy also is one of the core medical-ethics principles. In that context, autonomy dictates that patients should be able to make informed decisions about their care [Feil-Seifer and Mataric 2011]. This is reflected in the sample in which autonomy is discussed in relation to consent ($n = 11$), harm ($n = 9$), and health ($n = 7$).

Another issue broadly addressed is whether machines can have (a degree of) autonomy [Weng 2010]. In the reviewed literature, autonomy was also discussed rather frequently with robots ($n = 8$) and the Internet ($n = 12$). The connection between autonomy and the Internet was evident in connection to privacy issues of sharing data online.

As discussed earlier, autonomy also is connected with professionalism ($n = 11$). It is, for instance, argued that computing professionals have a duty to serve the public interest beyond the interests of employers or clients [Mitcham 2009].

5.1.4. Agency. The idea of agency is “conceptually associated with the idea of being capable of doing something that counts as an act or action” [Himma 2009]. Someone or something then is considered an agent if and only if capable of performing actions. Conceptually, agency has a close relationship to autonomy. In order to be considered a true agent, the capability of the agent must include the ability to make decisions underlying its actions, hence the agent must display some form of autonomy. This relationship between agency and autonomy was found in the literature ($n = 7$).

Of the 51 papers discussing agency, most are concerned with something other than someone being an agent. A common discussion concerns the extent to which intelligent artifacts such as robots can or should be considered agents [Floridi and Sanders 2004] or moral agents that can be held accountable for their behavior [Himma 2009].

Increasingly, autonomous artifacts as agents may create harm to human and nonhuman agents, or may suffer harm themselves and therefore should be considered moral agents [Sullins 2005]. This raises questions that range from whether machines can or should be moral decision-makers [Wallach 2008], whether artificial agents are morally responsible for their acts, and whether and how artificial (moral) agency will shift the burden for ethical behavior away from designers and users and onto the computer systems themselves [Allen et al. 2005].

This general trend was reflected in the technologies discussed most frequently in papers dealing with agency: artificial intelligence ($n = 11$), artificial agents ($n = 10$), and robots ($n = 16$). Additionally, a key ethical theory mentioned in relation to agency is Asimov’s laws of robotics ($n = 5$). Finally, the issues discussed most frequently in the sample in relation to agency are issues related to “specific computing technologies” ($n = 12$), which include issues specific to artificial intelligence and robotics.

5.1.5. Trust. Trust is a multifaceted concept often connected to privacy and security, chiefly concerning the processing of data. The reviewed literature reflected this theoretical overlap, with privacy and trust frequently discussed together ($n = 18$). Conceived as a component of privacy, trust is a characteristic of relationships between data authors, owners, data controllers, and analysts, in which there is mutual confidence in the appropriate use of information systems to process data in an acceptable manner. Trust is an evaluation of the perceived credibility, motivation, transparency, and responsibility of a system, its designers, and operators. Here, *credibility* concerns reputation [Little and Briggs 2009; Rashid et al. 2007, p. 190]; the system and its operators must be thought sufficiently *responsible* to handle one’s data well. *Motivation* regards the operators’ actual and intended uses of user data, the communication of which requires a degree of *transparency*. When a trusting relationship exists, data processing can proceed uninhibited within an agreed framework of accepted practices or purposes.

According to this approach, systems and operators with clearly defined hierarchies of responsibilities, in addition to frameworks to identify acceptable uses and oversight mechanisms for users, are increasingly trustworthy.

While traditionally a concept that describes the nature of the relationship between human actors mediated by information systems, trust can similarly exist between stakeholders and systems themselves [Li and Buchthal 2012], or as a critical component of human-computer interactions. In the reviewed literature, trust was often discussed in the context of the Internet ($n = 13$), wherein trusting relationships can be increasingly difficult to establish due to users entering into new relationships with each visit of a website. Similarly, users interact directly with an information system, meaning that the character of the system's operators may not be accessible [Mack 2004]. Similarly, users may unknowingly interact with, or be subjected to, the operations of automated "sorting" mechanisms, including algorithmic profiling [Schermer 2011; Mittelstadt and Floridi 2015] and surveillance technologies (see Lyon [2003]). Information systems augment human cognitive and decision-making capacities in processes involving normative judgments and tangible effects for users. In doing so, moral accountability traditionally assigned to human actors is displaced. Users may increasingly need to trust in information systems to process their inputs and data in acceptable ways, quite apart from how such data is eventually used by human actors. Without engaging in a lengthy discussion of the potential value neutrality and moral agency of computing technologies (e.g., Latour and Venn [2002]), it is recognized that, as computing increasingly plays a role in personal data processing and decision-making processes, identifying when and how trust can justifiably exist between users, systems, and operators is not a straightforward task.

5.1.6. Consent. Substantial attention has been given to issues of informed consent within research and commercial contexts ($n = 43$). Consent was often mentioned in connection to the Internet ($n = 12$) and medical technologies, including bio-ICT ($n = 6$) and e-health ($n = 4$). Informed consent is an ideal governance mechanism for research and development intended to prevent abuses in studies requiring voluntary participation [General Medical Council 2008; Angrist 2009]. Specified originally to govern medical research [Kious 2001], informed consent requires that users be provided sufficient information and guidance to understand the scope, purpose, and potential outcomes of a study, including known risks and benefits, the likelihood of each (including degrees of uncertainty), the identity of parties involved in designing and carrying out the study, any conflicts of interest, and participants' rights, including the right to withdraw. Consent procedures also tend to require in-built protections for participants, including mechanisms to prevent publication of identity, such as anonymization [Heeney 2012]. As information systems are increasingly developed and deployed by commercial entities beyond research settings subjected to ethical governance, consent also refers to agreements made between users and operations of a system defining acceptable uses and responsibilities. Comparable ethical oversight may be required in commercial contexts [Brown et al. 2004] due to potential harms being equivalent across academic and commercial usage.

Emerging computing technologies present numerous difficulties for informed consent. A key difficulty facing consent procedures for emerging computing technologies concerns defining technology-specific thresholds for informed consent (e.g., Kaplan and Litewka [2008]). The source of this problem is described by the dilemma of control [Collingridge 1980], which states that, as the implications of an emerging technology become known through its deployment, the difficulty of changing the technology's trajectory increases in parallel. At early stages of development and deployment, when modifying the technology would be comparatively simple, uncertainty remains over

its social and ethical effects, and thus the efficacy of proposed changes. Much of this uncertainty stems from potential uses of the data authored and processed [Mittelstadt and Floridi 2015], precluding the possibility of “well informed” consent³ in research and development of emerging computing technologies. At a minimum, single-instance models of consent [Choudhury et al. 2014] are increasingly ill suited to an “information age” marked by greater sharing of data across technological, social, commercial, and national borders [Clayton 2005; Ioannidis 2013; Mittelstadt and Floridi 2015]. Apart from dataflows, consent faces conceptual difficulties when applied to complex information systems, the workings (and thus implications) of which may not be transparent or comprehensible to users, acting as a limit on user autonomy ($n = 11$) and control over personal data (or privacy, $n = 33$).

5.1.7. Research Ethics. Research ethics is strongly linked to issues of consent ($n = 9$). Beyond issues of consent already specified, several papers identified emerging challenges for broader research ethics stemming from emerging capacities for Internet-based research ($n = 11$). Here, concerns with consent overlap with a broader focus on establishing norms of ethical conduct for research conducted in online and virtual environments [Varnhagen et al. 2005]. The need for new norms [Wilkinson and Thelwall 2011] stems from three facets of Internet research: (1) the temporal shifts in data retention and the accessibility of a “digital memory” implied by Internet technologies [Van Wel and Royakkers 2004; Oboler et al. 2012]; (2) the related opportunities to conduct research with unwilling participants using “scraped” data [Lomborg and Bechmann 2014; Markowetz et al. 2014; Krotoski 2012]; and (3) opportunities to engage participants anonymously and at a distance [Buchanan and Hvizdak 2009]. Concerning the first two, the Internet enables research without original data authorship through measurements, interviews, observations, and the like. Instead, online research can access extensive datasets retained in perpetuity across the Internet, seen, for example, through public forums and social media platforms. Similarly, access can be negotiated via platform operators rather than users (or participants), as much of what is available is considered public (e.g., amenable to research or third-party analysis) by default.

Such new avenues of research create new difficulties beyond issues of consent, concerning, for instance, ownership of data, replication of methods, recruitment and consent procedures for participations, distinguishing public and private digital spaces [Buchanan and Hvizdak 2009], how best to respect expectations of privacy of users in such public digital spaces, and how to anonymize publicly available data prior to the publication. Questions can also be raised over the capacities of existing ethics review boards to appreciate the unique challenges presented by online research [Zimmer 2010], and to define requirements striking an appropriate balance between potential contributions to knowledge and risks of research without explicit consent [Fairfield and Shtein 2014].

5.1.8. Identity. Identity can refer to two broad areas of enquiry: first, as in “identifiable information” concerning issues of data protection and anonymization (e.g., Mizani and Baykal [2007]); second, as a description of a person’s sense of self, which can be constrained in technologically mediated relationships (e.g., Mordini and Ottolini [2007]), for example, by constructing new informational representations of the user (e.g., Floridi [2011]). This differs from the conception in technical contexts, in which

³Here, a comparison can be made to consent procedures for medical treatments, for which a base of clinical evidence already exists, meaning that the likelihood and nature of risks can be communicated to patients during the consent process. Uncertainty here precludes meaningful communication of risk, beyond making clear that the risks of participation are largely unknown or unknowable at the time of consent.

identity is strongly linked with concepts such as authentication or credentials. In our analysis, these were captured under the category of security.

Concerning the concept of identity as identifiable information, forms of privacy that restrict the collection of data about the user are particularly relevant to enabling user control over identity, acting as barriers to “depersonalization” occurring when data representations of the user are constructed by third parties [Nagel and Remmers 2012]. Privacy as a description of a person’s sense of self is increasingly relevant as personal datasets are linked across platforms, both online and offline: for example, by linking the identity of users to online avatars [Girvan and Savage 2012].

In treating identity as an ethical concept, a right of control over (personal) identity is often advanced to ground normative claims. Papers discussing identity implicitly acknowledged this need for control, considering the importance of identity maintenance in the context of autonomy ($n = 7$), harm ($n = 9$), and privacy ($n = 16$). As an example, control over self-identity through management of information used to construct identity (e.g., Floridi [2011]) can be grounded by a desire to respect individual autonomy [Manders-Huits 2010]. Potential harms to an individual’s right to self-identity were linked to violations of privacy, wherein a lack of control over person data change those with access view and interact with the user. Interestingly, concerns over identity were not localized around any particular set of technologies, while being most commonly mentioned in relation to the Internet ($n = 6$), where users may increasingly be identified as “cyborgs” [Schultze and Mason 2012]. Human interaction with information systems can further transform identity in unexpected ways, for instance, by creating a sense of distance between actors that lessens the experience of moral agency, as has been observed in online contexts involving illegal or immoral practices [Chiou et al. 2012; Giustozzi and Van der Veer Martens 2011].

5.1.9. Inclusion, Digital Divides. The growing importance of computing as an enabler of economic activities, but also of social interaction, means that a lack of ability to participate in these activities can be a significant ethical concern [Hacker et al. 2009]. Questions of inclusion and divides arise within countries or societies as well as between them. Causes of digital divides are many and often reflect broader social divisions, such as those created by the economic system [Parayil 2005], or other types of inequality, such as those brought about by education, socioeconomic status, or age [Niehaves and Plattfaut 2014].

As digital divides relate to the individual and collective life chances of those affected by them, it is not surprising that, from the 45 papers in the sample that discuss inclusion and digital divides as issues, ethical theories related to justice ($n = 5$), human rights ($n = 4$), and Rawls, a philosopher well known for his views on justice as fairness ($n = 4$), were often employed. Inclusion and digital divides can raise ethical concerns with regard to any technology, but their problematic nature is most pronounced with regard to the Internet as a key enabling technology. It is therefore not surprising that the most frequently discussed technology with regard to inclusion issues was the generic category of ICT ($n = 14$), followed immediately by the Internet ($n = 13$).

The purpose of this overview of main issues was to give a quick indication of the type of problem discussed as ethical issues in computing and the relationship that they have to the other types of categories used in the survey. We focused on those categories that have a relatively well-defined content and are established discourses within ethics and computing, leaving aside more generic categories such as harm, misuse, deception, health-related issues, justice, behavior, practical issues, epistemological issues, design, education-related issues, moral values, and duties. This section has shown the content of a number of key issues, setting the scene for the overall discussion of the discourse on ethics and computing.

Fig. 14. Temporal development of Top 10 ethical issues.

5.2. The Discourse on Ethics and Computing

This article has analyzed the discourse on ethics and computing that is available in mainstream academic publications indexed in key databases to which researchers tend to have access. This discourse is important to understand for the technical community for its members to be sensitized to possible issues regarding the technologies they work with, while being provided with ways of addressing them.

This section discusses general trends and gaps in the overall discourse. A general observation is that the discourse is characterized by a very rich landscape of topics, in particular, many different types of ethical issues. We have been able to show only a high-level overview of this complex tapestry. The overall increase in number of publications in the area, as expressed in Figure 2, indicates a trend toward growing awareness that underlines the increasing importance of the field.

We assumed that there would be identifiable ethical issues or technologies whose importance became clear during the decade that we investigated. In order to test this assumption, we examined the temporal development of the top 10 ethical issues and technologies. Interestingly, no clear trend emerged from this analysis, as reflected in Figure 14, which shows temporal distribution for the top 10 ethical issues.

This figure shows that privacy is the dominant topic of ethics and computing and that this has not changed over the 10-year period under investigation. Somewhat surprisingly, the distribution of topics has been broadly constant, with no ethical issue emerging completely new or disappearing entirely, and the order being relatively static. A similar observation can be made with regard to the technologies described in the literature.

This finding is surprising, as we assumed that, in light of the rhetoric surrounding the disruptive nature of innovation in computing, there would be clearly identifiable trajectories of key ethical issues or technologies. Our data do not support this contention

Fig. 15. Quantity of papers coded across main categories (n = 599).

and show that the published literature on ethics and computing is rather static both in terms of ethical issues and computing technologies.

A further surprising finding was the limited use of ethical theory. Questions of ethics have been part of academic debate since the time of antiquity and have been discussed intensively over centuries. While much of this debate may not be relevant to computing questions, we assumed that it would be relatively standard to make explicit use of ethical theory. This does not imply that there would be agreement on which theory might be appropriate in a particular case, as philosophical ethics tends not to reach such final consensus [Russell 2001]. However, we assumed that a reference to ethical theory would clarify the argument. Contrary to our expectations, the majority of the papers in our sample fail to make use of ethical theory. A possible explanation of this may be the multidisciplinary nature of the community that contributes to the ethics of computing. While this community includes philosophers, who focus to a large degree on ethical theory, it also includes others for whom ethical theory does not seem to be relevant. An exacerbating factor may be that the term “ethics” is widely used and authors may feel justified in working with their intuitive understanding of the everyday use of the term.

The multidisciplinary nature of the contributors to the discourse may also explain why there is relatively little reference to methodology in the papers that we surveyed. As indicated earlier, explicit attention to methodology seems to be required in much of the social sciences, but it is much less prominent in other areas, including philosophy or computer science. This does not mean that members of these communities do not use a methodology, but it may imply that they do not feel a need to discuss it explicitly in such terms. In other words, certain methodologies are perceived as self-evident. For example, the author of a paper entirely discussing theoretical issues may assume that usage of a “conceptual” methodology is obvious to the reader. However, the existence of extensive background discourses on empirical and conceptual methodologies across the social sciences and philosophy suggest that methodologies should be made more explicit than is currently the case.

Some of the most interesting findings from our survey referred to these and other absences in the literature. As seen in Figure 15, the number of papers recorded as “none” (e.g., not given a substantive code) in the main categories reflects the degree

to which our initial assumption about the typical structure of the paper on ethics and computing was accurate.

Figure 15 also reveals that the assumption that papers would be likely to specify at least one ethical issue and make reference to a technology was confirmed by the vast majority of papers. Ethical theory and methodology played a much less prominent role, with only one-third to one-half of papers explicitly mentioning them.

The final two categories (recommendation and contribution) shed light on the internal understanding of the discourse by its authors and contributors. It provides interesting insights into why authors contribute and what they believe the consequences of their publications should be. The fact that less than one-third of the papers explicitly refer to their perceived contribution suggests a belief among authors that this aspect is self-evident and not in need of detailed discussion. This, in turn, can be interpreted as a sign that the authors have a clear view of the audience and they believe that the audience is capable of determining the contribution of their work. In light of the multidisciplinary nature of the discourse on ethics of computing discussed earlier, this can be seen as a problematic assumption. It is not obvious that audiences from other disciplines will have an immediate appreciation of the contribution of a paper unless this contribution is explicitly spelled out.

We therefore submit that there is a structural problem in the discourse concerning the ethics of computing. If this discourse truly aims at identifying ethical issues in computing and providing ways of addressing them, then papers need to be structured so that the various contributing disciplines are able to access and act on them.

At the moment, the results of the survey suggest that this recommendation has not been put into action. Authors appear to pursue individual disciplinary aims. This assessment is supported by the final category: Recommendations. If better understanding of the issues and changes to practice are the aims of research on ethics and computing, then one could expect to find practical and actionable recommendations as a matter of course. This does not imply that every paper has to have a simple checklist, but it might help if papers explicitly stated what should follow from them. While more than half of the papers in our sample include recommendations, the majority of these are not practical, but rather call for more research and discussion. No doubt, these are valuable recommendations as better understanding and broader debate of ethics and computing is desirable. However, in many cases, more practical guidance for specific stakeholders or groups might help identify and address such issues. The percentage of papers providing such practical guidance is relatively low, pointing toward a need for more practical guidance as a bridge between the theoretical concerns of philosophical ethics and the practical work of engineers, developers, regulators, and other ICT professionals.

6. CONCLUSION: ALIGNING ETHICS OF COMPUTING AND RESPONSIBLE RESEARCH AND INNOVATION

The importance of ethical considerations in computing continues to grow due to the ever increasing reach of computing technologies and artifacts into all areas of personal and public life in industrialized societies. While most members of these societies can be affected by such ethical issues, a key role is played by members of the technical communities who research, design, develop, and implement these technologies.

This article presents a survey of the literature on ethics and computing as represented in mainstream academic journal publications between 2003–2012. It provides an overview of ethical issues, technologies, ethical theory, methodology, recommendations, and contributions of this discourse by presenting the results of a qualitative study of 599 papers published during this period. Focusing on this decade

allowed us to demonstrate the development of the discussion as a whole. Individual issues, technologies, methodologies, and the like may vary. However, the overview of the entire discourse provides lasting insights that can inform computer scientists with regard to their individual and professional responsibilities. Based on a general structure of a paper on ethics in computing, we could show that some aspects of the discourse are well developed. The survey shows that there is a very rich discussion of a broad range of ethical issues related to computing, as well as a detailed interest in numerous technologies. Ethical theory is broadly discussed, albeit in a less detailed form than the ethical issues it can be used to describe. Other aspects are much more nascent. Methodology, recommendations, and contribution of papers are not uniformly well developed.

This article shows that there are rich sources for members of technical communities to make use of to gain a detailed understanding of ethical issues in computing. The article further makes a contribution to the awareness of such ethical issues by describing key ethical issues and their links to other categories. At the same time, the article points to some of the shortcomings of current research on ethics and computing. While the societal relevance of computing and its ethical consequences call for a collaboration between experts in different fields and disciplines, the analysis of the papers indicates that many of the authors involved in researching the ethics of computing remain wedded to their disciplinary traditions and fail to provide actionable advice to relevant stakeholders. This observation is the basis of our recommendations, which are described following the discussion of the limitations of our approach.

6.1. Limitations

The research described in this article is based on a detailed understanding and qualitative analysis of a large number of texts. This means that it is based on a large number of judgment calls by the researchers when delineating categories or nodes. For example, we coded “privacy” as an ethical issue and subsumed questions of data protection under it that could be equally relevant to law, regulation, identity, or security. Many of our codes had a large degree of ambiguity and could fit into more than one category. To address this, we developed a rigorous process of cross-checks and peer reviews among the authors, which we believe led to a robust outcome. We concede, however, that different interpretations of the literature are possible.

A second limitation is linked to the need to summarize and systematize a large body of work in a very short space. We chose to use a number of quantitative measures and graphical representations in order to convey key messages. We realize that this might be misconstrued as an attempt to provide a positivist and objectivist reading of literature, which we do not believe to be possible, in particular given the qualitative and interpretive nature of the coding process.

Despite these limitations, we believe that this article achieved its purpose in representing the discourse on ethics and computing, showing its strengths and weaknesses, thus supporting our recommendations as outlined in the next section.

6.2. Implications and Recommendations: The Transition to Responsible Research and Innovation in ICT

The ethics of computing forms part of what is typically called “applied ethics.” However, as can be seen from the analysis presented earlier, it is often not very applied or even applicable. Our research has demonstrated that the discourse, despite rapidly changing technologies and applications, remains relatively static. Over the course of the first decade of the 21st century, issues, theories, and approaches in ethics and computing have not changed substantially. On the one hand, this may be a positive sign, showing

the continuity of research and a certain level of resistance to fashions and fads. On the other hand, it may mean that the ethics of computing as an academic discourse is somewhat decoupled from the technical developments that it aims to reflect on.

If, as we believe, work on ethics and computing is to have practical relevance and influence technology research and design as well as policy and practice, then the field will have to develop substantially. We believe that, in addition to exploring and defining ethical issues, there should be more consistent attention to underlying ethical theory as well as practical implications, recommendations, or guidelines that help individuals deal with the ethical issues described earlier. To this end, we believe that a more explicit discussion of the contribution of individual pieces of research, as well as the methodology employed to arrive at particular insights, would be helpful.

The criticism that research on ethics and computing lacks impact is not novel. Our research has demonstrated that there are good reasons to believe this to be a valid criticism. At the same time, one can observe developments in the research landscape that are likely to alleviate this problem. At present, there is growing attention to the novel concept of responsible research and innovation (RRI) that aims to affect research and research governance with a view to having a practical impact on the way in which research and innovation activities relate to societal concerns.

ICT is not the starting point of the debate on RRI, which originates from technologies that are more obviously of ethical relevance, such as nanotechnology [Grunwald 2010], synthetic biology [Gutmann 2011; Zhang et al. 2011] or geoengineering [Macnaghten and Owen 2011]. However, there is increasingly a recognition that the ethical issues related to ICT require their own forms of responsible stewardship and management [Von Schomberg 2011].

A popular concept of RRI that has been adopted by the UK Engineering and Physical Sciences Research Council (EPSRC)⁴ was developed by Stilgoe et al. [2013] and is based on the four principles of Anticipation, Engagement, Reflection, and Action (AREA). To put it differently, research and innovation activities should aim to take into consideration likely and foreseeable consequences of their action, engage with a range of stakeholders likely to be affected by these activities, and reflect on purpose and underlying assumptions, all of which should lead to manifest actions that can change the course of the underlying research and innovation activities.

We believe that the discourse on ethics and computing, as described in this article, does a good job with regard to two of these points: anticipation and reflection. Reading the literature, one can get a detailed understanding of current ethical issues but also of likely future consequences of current technology development. The point of reflection is also well covered in the discourse, which discusses assumptions and underpinnings of research and questions of the purpose of research. The third aspect of RRI, the question of engagement, has a long tradition in computing research with regard to the inclusion of potential users in research and design. What is less prominent at present is a wider societal engagement and broader discussions not only about particular artifacts but about entire research agendas. This leads to a recommendation to research policymakers to encourage wider societal debate on research agendas.

The final point of RRI that we believe to be crucial to the conclusion of this article is the question of action. This refers to the question of how research on ethics and computing can have relevant and manifest consequences. We have argued earlier that, at present, the work on the ethics of computing is unlikely to be practically relevant. We have furthermore argued that more attention to the following points in research related to the ethics of computing is likely to overcome this issue to some degree:

⁴<https://www.epsrc.ac.uk/research/framework/>.

- Contribution to knowledge of the research
- Practical recommendations
- Transparent methodology

We assume that more specific research on particular technologies or ethical issues is more likely to have practical consequences than broader work on ICTs or computing in general. We realize that this focus on changing the nature of research on ethics and computing will not suffice. Scholars tend to follow their incentives. As long as incentives do not include a broader relevance of research, scholars are unlikely to change their focus. There are numerous initiatives that are trying to change the exclusive attention to the reception of publications by the peers of the scholars. This includes the requirement of European research proposals to discuss project impact or the UK Research Excellence Framework and its inclusion of impact in the evaluation criteria. The impact agenda that proposes to evaluate broader societal relevance of research as part of the overall understanding of science and innovation is highly contentious. From the point of view of ethics of computing, we believe nevertheless that it has the potential to incentivize researchers to pay more attention to ethical and social issues. Such a development will require a broad array of changes, starting from a changing focus of policymakers to different incentive structures on the broader scientific and organizational level, all the way to individual awareness of the concerned scholar. It implies that there will be culture changes that move ethics closer to criteria of scientific excellence as well as support structures that will allow research institutions and individuals to learn about ethics and ways of recognizing and dealing with it.

Computer scientists and technical experts play a key role in this context. Their understanding of technical principles and artifacts is key to understanding the ethical qualities of computing. The survey of the literature presented in this article should provide them with access to the discussion and a high-level overview of the subject. A single paper does not offer the space to provide detailed guidance or examples of all issues or technologies, but we believe that it has drawn out the main themes and given an overview of the link between computing and ethics. Understanding this complex relationship will help them to comply with the demands arising from their role as professionals. It will furthermore allow them to drill deeper in particular areas as they reflect on the potential impact of current research or practice. Finally, we hope that it will motivate them to seek collaboration with experts in ethical and social aspects, as a more comprehensive understanding of the ethics of computing and resulting changes to practice will require closer collaboration of experts from these different fields.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 609817 (Responsible-Industry; www.responsible-industry.eu) and no. 604102 (Human Brain Project). The research furthermore benefited from the activities undertaken in the project "Framework for Responsible Research and Innovation in ICT," EPSRC reference EP/J000019/1. The authors would like to thank Helge Janicke for helpful comments on an earlier version and the anonymous reviewers of the *CSUR* journal for insights into how to improve the paper.

REFERENCES

- Alison Adam. 2001. Computer ethics in a different voice. *Information and Organization* 11, 4, 235–261.
- M. Al Ameen, J. Liu, and K. Kwak. 2012. Security and privacy issues in wireless sensor networks for healthcare applications. *Journal of Medical Systems* 36, 1, 93–101.
- Bill Albrecht, Ken Christensen, Venu Dasigi, Jim Huggins, and Jody Paul. 2012. The pledge of the computing professional: Recognizing and promoting ethics in the computing professions. *SIGCAS Computers and Society* 42, 1, 6–8. DOI: <http://dx.doi.org/10.1145/2422512.2422513>

- C. Allen, I. Smit, and W. Wallach. 2005. Artificial morality: Top-down, bottom-up, and hybrid approaches. *Ethics and Information Technology* 7, 3, 149–155.
- P. J. Allen and L. D. Roberts. 2010. The ethics of outsourcing online survey research. *International Journal of Technoethics* 1, 3, 35–48. DOI: <http://dx.doi.org/10.4018/jte.2010070104>
- F. Allhoff, P. Lin, J. Moor, and J. Weckert. 2010. Ethics of human enhancement: 25 questions & answers. *Studies in Ethics, Law, and Technology* 4, 1. DOI: <http://dx.doi.org/10.2202/1941-6008.1110>
- M. L. Anderson. 2005. Why is AI so scary? *Artificial Intelligence* 169, 2, 201–208.
- M. Angrist. 2009. Eyes wide open: The personal genome project, citizen science and veracity in informed consent. *Personalized Medicine* 6, 691–699.
- Aristotle. 2007. *The Nicomachean Ethics*. Filiquarian Publishing, LLC, Minneapolis, MN.
- S. A. Bagüés, A. Zeidler, C. Klein, F. Valdivielso, and R. Matias. 2010. Enabling personal privacy for pervasive computing environments. *Journal of Universal Computer Science* 16, 341–371.
- D. Banerjee, T. P. Cronan, and T. W. Jones. 1998. Modeling IT ethics: A study in situational ethics. *MIS Quarterly* 22, 1, 31–60.
- Jeremy Bentham. 1789. *An Introduction to the Principles of Morals and Legislation*. Dover Publications Inc., Mineola, NY.
- Sebastian Boell and Dubravka Cezec-Kecmanovic. 2014. A hermeneutic approach for conducting literature reviews and literature searches. *Communications of the Association for Information Systems* 34, 1.
- K. W. Bowyer. 2004. Face recognition technology: Security versus privacy. *IEEE Technology and Society Magazine* 23, 1, 9–20. DOI: <http://dx.doi.org/10.1109/MTAS.2004.1273467>
- P. Brey. 2005. Freedom and privacy in ambient intelligence. *Ethics and Information Technology* 7, 3, 157–166.
- Philip Brey and Johnny H. Soraker. 2009. Philosophy of computing and information technology. In Dov M. Gabbay, Anthonie W. M. Meijers, John Woods, and Paul Thagard (Eds). *Philosophy of Technology and Engineering Sciences: 9*. North Holland, 1341–1408.
- S. Brown, N. Hine, A. Sixsmith, and P. Garner. 2004. Care in the community. *BT Technology Journal* 22, 3, 56–64.
- Elizabeth A. Buchanan and Erin E. Hvizdak. 2009. Online survey tools: Ethical and methodological concerns of human research ethics committees. *Journal of Empirical Research on Human Research Ethics* 4, 2, 37–48.
- T. Bynum. 2008a. Norbert Wiener and the rise of information ethics. In Jeroen van den Hoven and John Weckert (Eds). *Information Technology and Moral Philosophy*. Cambridge University Press, New York, NY, 8–25.
- Terrell Bynum. 2008b. Computer and information ethics. In *Stanford Encyclopedia of Philosophy*. Retrieved January 14, 2016 from <http://plato.stanford.edu/entries/ethics-computer/>.
- Terrell Ward Bynum and Simon Rogerson. 2003. *Computer Ethics and Professional Responsibility: Introductory Text and Readings*. WileyBlackwell, Hoboken, NJ.
- C. Cary, H. J. Wen, and P. Mahatanankoon. 2003. Data mining: Consumer privacy, ethical policy, and systems development practices. *Human Systems Management* 22, 4, 157–168.
- W. B. Chiou, P. H. Wan, and C. S. Wan. 2012. A new look at software piracy: Soft lifting primes an inauthentic sense of self, prompting further unethical behavior. *International Journal of Human Computer Studies* 70, 2, 107–115.
- Suparna Choudhury, Jennifer R. Fishman, Michelle L. McGowan, and Eric T. Juengst. 2014. Big data, open science and the brain: lessons learned from genomics. *Frontiers in Human Neuroscience* 8, 239. DOI: <http://dx.doi.org/10.3389/fnhum.2014.00239>
- Ellen Wright Clayton. 2005. Informed consent and biobanks. *Journal of Law, Medicine & Ethics* 33, 1, 15–21. DOI: <http://dx.doi.org/10.1111/j.1748-720X.2005.tb00206.x>
- Eli Cohen and Larry Cornwell. 1989. A question of ethics: Developing information system ethics. *Journal of Business Ethics* 8, 6, 431–437.
- D. Collingridge. 1980. *The Social Control of Technology*. Palgrave Macmillan, Basingstoke, UK.
- Mary L. Cummings. 2006. Integrating ethics in design through the value-sensitive design approach. *Science and Engineering Ethics* 12, 4, 701–715. DOI: <http://dx.doi.org/10.1007/s11948-006-0065-0>
- J. K. Estell and K. Christensen. 2011. The need for a new graduation rite of passage. *Communications of the ACM* 54, 2, 113–115.
- Joshua Fairfield and Hannah Shtein. 2014. Big data, big problems: Emerging issues in the ethics of data science and journalism. *Journal of Mass Media Ethics* 29, 1, 38–51. DOI: <http://dx.doi.org/10.1080/08900523.2014.863126>

- D. Feil-Seifer and M. J. Mataric. 2011. Socially assistive robotics: Ethical issues related to technology. *IEEE Robotics and Automation Magazine* 18, 1, 24–31.
- Luciano Floridi. 2010. Information ethics. In Luciano Floridi (Ed.). *The Cambridge Handbook of Information and Computer Ethics*. Cambridge University Press, New York, NY, 77–97.
- Luciano Floridi. 2011. The informational nature of personal identity. *Minds and Machines* 21, 4, 549–566. DOI: <http://dx.doi.org/10.1007/s11023-011-9259-6>
- Luciano Floridi. 2013. *The Philosophy of Information* Reprint edition., Oxford University Press, Oxford, UK.
- Luciano Floridi and J. W. Sanders. 2004. On the morality of artificial agents. *Minds and Machines* 14, 3, 349–379. DOI: <http://dx.doi.org/10.1023/B:MIND.0000035461.63578.9d>
- M. Friedewald, O. D. Costa, Y. Punie, P. Alahuhta, and S. Heinonen. 2005. Perspectives of ambient intelligence in the home environment. *Telematics and Informatics* 22, 3, 221–238.
- Batya Friedman, Peter Kahn, and Alan Borning. 2008. Value sensitive design and information systems. In Ken Himma and Herman Tavani (Eds.). *The Handbook of Information and Computer Ethics*. Wiley Blackwell, Hoboken, NJ, 69–102.
- General Medical Council. 2008. Consent: patients and doctors making decisions together. Retrieved January 14, 2016 from http://www.gmc-uk.org/guidance/ethical_guidance/consent_guidance_index.asp.
- Carol Gilligan. 1990. *In a Different Voice: Psychological Theory and Women's Development*. Harvard University Press, Cambridge, MA.
- C. Girvan and T. Savage. 2012. Ethical considerations for educational research in a virtual world. *Interactive Learning Environments* 20, 3, 239–251.
- Emilie Steele Giustozzi and Betsy Van derVeer Martens. 2011. The new competitive intelligence agents: “Programming” competitive intelligence ethics into corporate cultures. *Webology* 8, 2.
- M. E. Gorman. 2001. Turning students into ethical professionals. *IEEE Technology and Society Magazine* 20, 4, 21–27. DOI: <http://dx.doi.org/10.1109/44.974504>
- Don Gotterbarn, Keith Miller, and Simon Rogerson. 1999. Computer society and ACM approve software engineering code of ethics. *Computer* 32, 10, 84–88.
- Armin Grunwald. 2010. From speculative nanoethics to explorative philosophy of nanotechnology. *Nanoethics* 4, 2, 91–101. DOI: <http://dx.doi.org/10.1007/s11569-010-0088-5>
- Amy Gutmann. 2011. The ethics of synthetic biology: Guiding principles for emerging technologies. *Hastings Center Report* 41, 4, 17–22.
- K. L. Hacker, S. M. Mason, and E. L. Morgan. 2009. Digital disempowerment in a network society. *International Journal of Electronic Government Research* 5, 2, 57–71.
- C. Heeney. 2012. Breaching the contract? Privacy and the UK census. *Information Society* 28, 5, 316–328.
- E. Hildt. 2010. Brain–Computer interaction and medical access to the brain: Individual, social and ethical implications. *Studies in Ethics, Law, and Technology* 4, 3.
- K. E. Himma. 2009. Artificial agency, consciousness, and the criteria for moral agency: What properties must an artificial agent have to be a moral agent? *Ethics and Information Technology* 11, 1, 19–29.
- K. E. Himma and H. T. Tavani. 2008. *The handbook of information and computer ethics*. Wiley, Hoboken, NJ.
- Jeroen van den Hoven and John Weckert (Eds.) 2008. *Information Technology and Moral Philosophy*. Cambridge University Press, New York, NY.
- Gordon Hull. 2012. Robert Merges: Justifying intellectual property. *Ethics and Information Technology* 14, 2, 169–177. DOI: <http://dx.doi.org/10.1007/s10676-012-9287-9>
- John P. A. Ioannidis. 2013. Informed consent, big data, and the oxymoron of research that is not research. *American Journal of Bioethics* 13, 4, 40–42. DOI: <http://dx.doi.org/10.1080/15265161.2013.768864>
- Deborah G. Johnson. 1985. *Computer Ethics*. Prentice-Hall, Upper Saddle River, NJ.
- Deborah G. Johnson and Helen Nissenbaum. 1995. *Computers, Ethics and Social Values*. Pearson, New York, NY.
- Damien Joseph, Kok-Yee Ng, Christine Koh, and Soon Ang. 2007. Turnover of information technology professionals: a narrative review, meta-analytic structural equation modeling, and model development. *MIS Quarterly* 31, 3, 547–577.
- Immanuel Kant. 1797. *Grundlegung zur Metaphysik der Sitten*. Reclam, Ditzingen.
- Immanuel Kant. 1788. *Kritik der praktischen Vernunft*. Reclam, Ditzingen.
- Bonnie Kaplan and Sergio Litewka. 2008. Ethical challenges of telemedicine and telehealth. *Cambridge Quarterly of Healthcare Ethics* 17, 4, 401–416.
- B. M. Kious. 2001. The Nuremberg Code: its history and implications. *Princeton Journal of Bioethics* 4, 7–19.

- Aleksandra K. Krotoski. 2012. Data-driven research: Open data opportunities for growing knowledge, and ethical issues that arise. *Insights: the UKSG Journal* 25, 1, 28–32. DOI: <http://dx.doi.org/10.1629/2048-7754.25.1.28>
- V. Lally, M. Sharples, F. Tracy, N. Bertram, and S. Masters. 2012. Researching the ethical dimensions of mobile, ubiquitous and immersive technology enhanced learning (MUITEL): A thematic review and dialogue. *Interactive Learning Environments* 20, 3, 217–238.
- Bruno Latour and Couze Venn. 2002. Morality and technology: The end of the means. *Theory Culture Society* 19, 5–6, 247–260. DOI: <http://dx.doi.org/10.1177/026327602761899246>
- M. Li and S. Buchthal. 2012. Advisory services in the virtual world: An empowerment perspective. *Electronic Commerce Research* 12, 1, 53–96.
- L. Little and P. Briggs. 2009. Pervasive healthcare: The elderly perspective. In *2nd International Conference on Pervasive Technologies Related to Assistive Environments (PETRA'09)*. Corfu.
- Stine Lomborg and Anja Bechmann. 2014. Using APIs for data collection on social media. *Inf. Soc.* 30, 4 (2014), 256–265. DOI: <http://dx.doi.org/10.1080/01972243.2014.915276>
- David Lyon. 2003. *Surveillance as social sorting: Privacy, Risk, and Digital Discrimination*. Routledge, London, UK.
- Alasdair C. MacIntyre. 2007. *After Virtue: A Study in Moral Theory*. University of Notre Dame Press, Notre Dame, IN.
- John Mack. 2004. Beyond HIPAA: Ethics in the e-health arena. Ethical guidelines help ensure that credible health information appears on the Internet. *Healthcare Executive* 19, 5, 32–3.
- Phil Macnaghten and Richard Owen. 2011. Good governance for geoengineering. *Nature* 479, 7373, 293. DOI: <http://dx.doi.org/10.1038/479293a>
- N. Manders-Huits. 2010. Practical versus moral identities in identity management. *Ethics and Information Technology* 12, 1, 43–55.
- Alexander Markowetz, Konrad Błaskiewicz, Christian Montag, Christina Switala, and Thomas E. Schlaepfer. 2014. Psycho-informatics: Big data shaping modern psychometrics. *Medical Hypotheses* 82, 4, 405–411. DOI: <http://dx.doi.org/10.1016/j.mehy.2013.11.030>
- M. Lynne Markus and Kevin Mentzer. 2014. Foresight for a responsible future with ICT. *Information System Frontiers* 1–16. DOI: <http://dx.doi.org/10.1007/s10796-013-9479-9>
- Maris Martinsons, et al. 2009. The Ethics of IT Professionals in Japan and China. *Journal of the Association for Information Systems* 10, 11, 834–859.
- A. Marturano. 2002. The role of metaethics and the future of computer ethics. *Ethics and Information Technology* 4, 1, 71–78.
- M. T. McFall. 2012. Real character-friends: Aristotelian friendship, living together, and technology. *Ethics and Information Technology* 14, 3, 221–230.
- K. W. Miller. 2009. Joining the conversation about IT ethics. *IT Professional* 11, 5, 48–50.
- John Stuart Mill. 1861. *Utilitarianism* (2nd rev. ed.), Hackett Publishing Co, Inc., Cambridge, MA.
- John Mingers. 2001. Combining IS research methods: Towards a pluralist methodology. *Information Systems Research* 12, 3, 240–259. DOI: <http://dx.doi.org/Article>
- John Mingers and Geoff Walsham. 2010. Towards ethical information systems: The contribution of discourse ethics. *MIS Quarterly* 34, 4, 833–854.
- C. Mitcham. 2009. Convivial software: An end-user perspective on free and open source software. *Ethics and Information Technology* 11, 4, 299–310.
- Brent Mittelstadt and Luciano Floridi. 2015. The ethics of big data: Current and foreseeable issues in biomedical contexts. *Science and Engineering Ethics* online firsts (2015). DOI: <http://dx.doi.org/10.1007/s11948-015-9652-2>
- M. A. Mizani and N. Baykal. 2007. A software platform to analyse the ethical issues of electronic patient privacy policy: the S3P example. *Journal of Medical Ethics* 33, 12, 695–698.
- Emilio Mordini and Corinna Ottolini. 2007. Body identification, biometrics and medicine: Ethical and social considerations. *Annali dell'Istituto Superiore di Sanita* 43, 1, 51–60.
- Arash Mostaghimi and Bradley H. Crotty. 2011. Professionalism in the digital age. *Annals of Internal Medicine* 154, 8, 560–W197.
- Saskia K. Nagel and Hartmut Remmers. 2012. Self-perception and self-determination in surveillance conditions. *American Journal of Bioethics* 12, 9, 53–55.
- Björn Niehaves and Ralf Plattfaut. 2014. Internet adoption by the elderly: employing IS technology acceptance theories for understanding the age-related digital divide. *European Journal of Information Systems* 23, 6, 708–726. DOI: <http://dx.doi.org/10.1057/ejis.2013.19>

- Niels Nijsingh and Marcus Duwell. 2009. Interdisciplinary, applied ethics and social science. In Paul Sollie and Marcus Düwell (Eds.). *Evaluating New Technologies: Methodological Problems for the Ethical Assessment of Technology Developments*. The International Library of Ethics, Law and Technology. Springer, New York, NY, 79–92.
- Helen Nissenbaum. 2004. *Privacy as Contextual Integrity*. Social Science Research Network, Rochester, NY.
- Andre Oboler, K. Welsh, and L. Cruz. 2012. The danger of big data: Social media as computational social science. *First Monday* 17, 7.
- Govindan Parayil. 2005. The digital divide and increasing returns: Contradictions of informational capitalism. *The Information Society* 21, 1, 41–51.
- Dinah Payne and Brett J. L. Landry. 2006. A uniform code of ethics: Business and IT professional ethics. *Communications of the ACM* 49, 11, 80–84.
- J. A. Quilici-Gonzalez, G. Kobayashi, M. C. Broens, and M. E. Q. Gonzalez. 2010. Ubiquitous computing: Any ethical implications? *International Journal of Technoethics* 1, 3, 11–23. DOI: <http://dx.doi.org/10.4018/jte.2010070102>
- U. Rashid, H. Schmidtke, and N. Woo. 2007. Managing disclosure of personal health information in smart home healthcare. In *4th International Conference on Universal Access in Human-Computer Interaction (UAHCI'07)*. Beijing.
- Frantz Rowe. 2014. What literature review is not: diversity, boundaries and recommendations. *European Journal of Information Systems* 23, 3, 241–255. DOI: <http://dx.doi.org/http://dx.doi.org.proxy.library.dmu.ac.uk/10.1057/ejis.2014.7>
- Bertrand Russell. 2001. *The Problems of Philosophy*. Oxford University Press, New York, NY.
- Fariba Sadri. 2011. Ambient intelligence: A survey. *ACM Computing Surveys* 43, 4, 36.1–36.66. DOI: <http://dx.doi.org/10.1145/1978802.1978815>
- P. Salvini, E. Datteri, C. Laschi, and P. Dario. 2008. Scientific models and ethical issues in hybrid bionic systems research. *AI and Society* 22, 3, 431–448. DOI: <http://dx.doi.org/10.1007/s00146-007-0158-6>
- Suprateek Sarker, Mark Fuller, and Sutirtha Chatterjee. 2009. Ethical information systems development: A Baumanian postmodernist perspective. *Journal of the Association for Information Systems* 10, 11, 787–815.
- Geoff Sayre-McCord. 2014. Metaethics. In Edward N. Zalta (Ed.). *The Stanford Encyclopedia of Philosophy*. Retrieved January 14, 2016 from <http://plato.stanford.edu/entries/metaethics/>.
- B. W. Schermer. 2011. The limits of privacy in automated profiling and data mining. *Computer Law and Security Review* 27, 1, 45–52.
- Ferdinand Schoeman. 1984. Privacy and intimate information. In *Philosophical dimensions of privacy: An anthology*, Ferdinand Schoeman (Ed.). Cambridge University Press, Cambridge. 403–408.
- U. Schultze and R. O. Mason. 2012. Studying cyborgs: Re-examining internet studies as human subjects research. *Journal of Information Technology* 27, 4, 301–312.
- Richard A. Spinello and Herman T. Tavani. 2004. *Intellectual Property Rights in a Networked World: Theory and Practice*. Information Science Publishing, Hershey, PA.
- Richard A. Spinello and Herman T. Tavani. 2001. *Readings in CyberEthics*. Jones and Bartlett Publishers, Inc., Burlington, MA.
- Bernd Stahl. 2012. Morality, Ethics, and Reflection: A categorization of normative IS research. *Journal of the Association for Information Systems* 13, 8.
- Jack Stilgoe, Richard Owen, and Phil Macnaghten. 2013. Developing a framework for responsible innovation. *Research Policy* 42, 9, 1568–1580. DOI: <http://dx.doi.org/10.1016/j.respol.2013.05.008>
- E. Stip and V. Rialle. 2005. Environmental cognitive remediation in schizophrenia: Ethical implications of “smart home” technology. *Canadian Journal of Psychiatry-Revue Canadienne De Psychiatrie* 50, 5, 281–291.
- J. P. Sullins. 2005. Ethics and artificial life: From modeling to moral agents. *Ethics and Information Technology* 7, 3, 139–148.
- David R. Tranfield, David Denyer, and Palminder Smart. 2003. Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management* 14, 207–222.
- J. van den Hoven. 2008. Information technology, privacy, and the protection of personal data. In *Information Technology and Moral Philosophy*. Cambridge Studies in Philosophy and Public Policy. Cambridge University Press, New York, NY.
- L. Van Wel and L. Royakkers. 2004. Ethical issues in web data mining. *Ethics and Information Technology* 6, 2, 129–140.

- C. K. Varnhagen et al. 2005. How informed is online informed consent? *Ethics & Behavior* 15, 1, 37–48.
- Rene Von Schomberg (Ed.): 2011. *Towards Responsible Research and Innovation in the Information and Communication Technologies and Security Technologies Fields*. Publication Office of the European Union, Luxembourg.
- Ken Wakita and Toshiyuki Tsurumi. 2007. Finding community structure in mega-scale social networks: [Extended Abstract]. In *Proceedings of the 16th International Conference on World Wide Web (WWW'07)*. ACM, New York, NY, 1275–1276. DOI: <http://dx.doi.org/10.1145/1242572.1242805>
- W. Wallach. 2008. Implementing moral decision making faculties in computers and robots. *AI and Society* 22, 4, 463–475.
- D. M. Wasieleski and M. Gal-Or. 2008. An enquiry into the ethical efficacy of the use of radio frequency identification technology. *Ethics and Information Technology* 10, 1, 27–40.
- John Weckert and Richard Lucas (Eds.). 2013. *Professionalism in the Information and Communication Technology Industry*. The Australian National University, Canberra, Australia.
- Joseph Weizenbaum. 1977. *Computer Power and Human Reason: From Judgement to Calculation*. W. H. Freeman & Co Ltd.
- Y. H. Weng. 2010. Beyond robot ethics: On a legislative consortium for social robotics. *Advanced Robotics* 24, 13, 1919–1926.
- Alan F. Westin. 1970. *Privacy and Freedom*. The Bodley Head, London, UK.
- Norbert Wiener. 1964. *God and Golem, Inc. A Comment on Certain Points where Cybernetics Impinges on Religion*. MIT Press, Cambridge, MA.
- Norbert Wiener. 1954. *The Human Use of Human Beings*. Doubleday, New York, NY.
- David Wilkinson and Mike Thelwall. 2011. Researching personal information on the public web methods and ethics. *Social Science Computer Review* 29, 4, 387–401. DOI: <http://dx.doi.org/10.1177/0894439310378979>
- T. J. Williamson. 2010. Predicting building performance: The ethics of computer simulation. *Building Research and Information* 38, 4, 401–410. DOI: <http://dx.doi.org/10.1080/09613218.2010.481204>
- D. Wright, S. Gutwirth, M. Friedewald, E. Vildjiounaite, and Y. Punie. 2008. *Safeguards in a World of Ambient Intelligence*. Springer-Verlag, New York, NY.
- Aimee van Wynsberghe. 2013. Designing robots for care: Care centered value-sensitive design. *Science and Engineering Ethics* 19, 2, 407–433.
- Joy Y. Zhang, Claire Marris, and Nikolas Rose. 2011. *The transnational governance of synthetic biology—Scientific uncertainty, cross-borderness and the “art” of governance*. London School of Economics and Political Science, London, UK.
- M. Zimmer. 2010. “But the data is already public”: On the ethics of research in Facebook. *Ethics and Information Technology* 12, 4, 313–325.

Received June 2015; revised November 2015; accepted November 2015