Workshop Ethics and Morality in Business Informatics Workshop Ethik und Moral in der Wirtschaftsinformatik (EMoWI 2020)

Jens Gulden[1], Alexander Bock[2], Sergio España[1]

[1] Utrecht University, Department of Information and Computing Sciences, Utrecht,
The Netherlands
{j.gulden,s.espana}@uu.nl
[2] University of Duisburg-Essen, Research Group Enterprise Modeling and Information
Systems, Essen, Germany
alexander.bock@uni-due.de

Abstract. With its second edition, the workshop on Ethics and Morality in Business Informatics (EMoWI) 2020 continues its project of providing a forum for research and debate on a severely underrepresented area of study in the field of Business Information Systems (BIS). The workshop series addresses itself to ethical issues arising from the development and use of BIS, as well as to ethical questions and conflicts residing in BIS research itself. The contributions of this year's workshop underline both the significance and the variety of research topics within this domain. A common methodological theme of these contributions is the use of traditional BIS approaches, such as conceptual domain analysis and method engineering techniques, to promote ethical reflections concerning both the subjects and the conduct of BIS research. This editorial provides an overview of the background and the contributions of the EMoWI workshop 2020.

Keywords: Ethics, Values, Morals, Moral Philosophy, Business Informatics, Business Information Systems.

1   Introduction

The impact of information technology (IT) on almost all areas of our lives has become strikingly obvious. IT developments of the past two decades have resulted in once unimaginable artifacts and possibilities, overthrowing conditions of human life that prevailed for centuries. The worldwide connection of humans through modern communication and social media technologies has led to the emergence of new types of social structures. Artificial intelligence (AI) technologies have eventuated in the unprecedented possibility to obtain knowledge claims not made by humans. These and other technological developments affect social, organizational, and societal processes in profound ways, and consequently they also give rise to novel ethical problems and questions. For example, the availability of AI-based knowledge claims raises the question of whether such knowledge claims may be used in critical decisions, such as life-or-death decisions. And the possibility to support the emergence of new social structures via communication platforms goes hand in hand with the possibility to suppress them. The present workshop is intended to provide a forum to discuss novel ethical problems and questions such as these.

One might argue that ethical questions are better left to the philosopher. We, however, think that there are good reasons why there is an important place for ethical considerations in the field of BIS research – and a more significant place than has hitherto been acknowledged. First of all, as just pointed out, many new ethical issues are the immediate consequence of new IT artifacts. We as Business Information Systems (BIS) researchers are among the prime drivers of socio-technological innovation. In particular, we are concerned with the development and use of IT artifacts in social and organizational context, and we provide tools and methods to develop information systems. Thus, one may argue that we have a responsibility for the ethical implications of the tools we study and whose development we support. Conversely, we as BIS researchers also bring special competencies to the study of ethical issues of information systems. Unlike other investigators, we are familiar with the internal functioning of IT artifacts, while also being acutely aware of their social and organizational ramifications. Thus, one may expect that BIS research can develop more nuanced accounts of ethical issues of information systems than other disciplines can.

It is, therefore, unfortunate that the study of ethical issues of BIS has not yet received the attention it deserves in our field. With the present workshop series, we wish to contribute to establishing the study of moral and ethical issues within BIS research. Following Kant's characterization of ethics as being concerned with the question "What ought I to do?", as opposed to the other two fundamental philosophical questions, "What can I know?" and "What may I hope?" [1, p. A805/B833], we may characterize ethics in BIS as the concern with the question, "What ought I to do as a Business Information Systems researcher?"

2 Contributions of the Workshop

All contributions to this year's EMoWI workshop, in some way or other, raise awareness of the need to supplement methods for system development and use with components to foster ethical reflections. Two contributions to the workshop directly address the incorporation of ethics into systems a development approach, namely *More than Ticking Off a Checklist?...* (Jahn et al.) and *A Research Commentary – Integrating Ethical Issues into the Data Process* (Levina). Two other contributions take in a meta-perspective. Conceptual reflection on the role of transparency for achieving accountability are provided by the contribution *Transparenz als Mittel der Verantwortlichkeit bei KI-gestützten Systemen* (*Transparency as means for accountability in AI-based systems*, Richter), and a general methodological perspective is taken by the contribution *Ethical Implications of Security Vulnerability Research for Critical Infrastructure Protection* (Obiora Nweke and Wolthusen)*,* relating traditional philosophical doctrines to research on IT security.

The paper *More than Ticking Off a Checklist? Towards an Approach for Quantifying the Effectiveness of Responsible Innovation in the Design Process* by Katharina Jahn et al. examines the impact of ethical considerations in system design. The fundamental problem is that the effectiveness of ethical considerations is difficult to evaluate empirically. This would require designing the same system at least twice, implement both variants, and then experimentally evaluate the systems' conformance to ethical values. It is obvious that such an experimental setup cannot be carried out with reasonable effort. To tackle this problem, the paper introduces a three-level approach to reflect on design decisions and their ethical implications. The first level represents an initial design decision without additional ethical reflection; the second level describes a modified decision meant to take place after a dedicated discussion among stakeholders, and the third level, possibly the best solution, corresponds to what Kant has introduced as a "regulative idea" [1, p. A644/B672]. This conceptual instrumentation enables to establish a qualitative distance measure between different design decisions, as it allows to locate the degree of ethical reflection on a three-fold described continuum. With such a conceptual distance measure at hand, different system designs can be systematically compared regarding the effectiveness of ethical reflection in design. This contributes to the idea of systematically incorporating ethics into design processes in BIS by reflexively applying the methodical apparatus of our disciplines to our methods.

The paper *A Research Commentary - Integrating Ethical Issues into the Data Process* by Olga Levina discusses how data-based software applications are subject to ethical concerns. Such applications are often central in data science projects but are also used for conventional organizational management, when embedded into BIS. Taking the data science process presented in [2] as a reference process, the paper identifies the stakeholder groups (roles) involved in it, and the activities it comprises, which are structured in phases (namely, Sense, Transform, Act and Apply) and stages. This process is then used as a framework to identify and discuss ethical issues that are

relevant to each phase and stage. Most issues refer to the social dimension, but there is also space for environmental concerns. For instance, the Sense phase, when data is collected, detected in the system and pre-selected for detailed analysis, is subject to data privacy issues (often regulated by existing laws), environmental aspects such as the power used by the storage technology, and often needs a careful discussion about the ownership of the data. During the Apply phase, when the software application is applied in practice and can thus impact the social or natural environments, several stakeholder groups might have concerns related to the purpose with which the application is used for and the effects of such use. This is especially critical in domains such as healthcare, the legal system, or the job market. The analysis also yields a set of potential questions to be asked to the stakeholders, so as to facilitate the discussion of ethical issues during the data science process and the proper management of countermeasures to reduce negative impacts. The paper opens more questions than it answers, becoming an interesting departing point for further research where computer and information scientists could and should collaborate with social scientists, ethicists, or philosophers, as well as with experts in the domains of specific applied data science projects.

Andreas Richter's paper *Transparenz als Mittel der Verantwortlichkeit bei KI-gestützten Systemen (Transparency as means for accountability in AI-based systems)* deals with fundamental issues of data exchange among parties of unequal power. In particular, attention is directed at the AI-based use of personal data by large corporations. Richter's principal means and focus of analysis is the concept of transparency, expanding on arguments previously developed by Ananny and Crawford (2018) [3]. According to these arguments, it is problematic to view transparency as an ideal in itself, and to assume that attaining transparency is tantamount to avoiding the dangers of data (mis-)use. Rather, it is argued, one needs to answer various questions pertaining to the context and interpretation of transparency. Relying on the previous analysis by Ananny and Crawford (2018), Richter first identifies four areas that need to be considered in examining transparency. He illustrates this classification by the example of data analysis for autonomous driving, pointing out that not all sorts of transparency are to the benefit of individual persons. Following this, the paper presents two thought experiments, further exemplifying how subtle differences in how data is shared and used can lead to very different judgments about whether doing so seems desirable. Ultimately, the paper demonstrates that careful distinctions are needed in evaluating the merits and dangers of data exchange in today's economy.

Fundamental difficulties BIS researchers face when confronted with ethical dilemmas are discussed in Livinus Obiora Nweke and Stephen D. Wolthusen's paper *Ethical Implications of Security Vulnerability Research for Critical Infrastructure Protection*. Along the traditional triad of deontological, consequentialist and virtue ethics, the paper exemplifies how ethical questions are addressed differently, depending on the methodical apparatus applied. At the core of the paper lies the scenario of a moral dilemma in which a security vulnerability researcher is involved after discovering a security risk in a critical infrastructure operated by an infrastructure provider. The paper constructs a thought experiment in which the only options to act are either to inform

the infrastructure provider about the security risk, or to inform the public. Doing the first would protect the provider's interest in reputation, but might tempt him to delay the mitigation of the risk, while doing the second would force the provider to act quickly, however, would expose the security risk to a greater danger of being exploited because it gets known to a large number of people. The paper subsequently shows that the three normative approaches in ethics lead to different resolutions of the dilemma. This discussion contributes to an understanding of ethics not as a methodical apparatus for giving definite answers to ethical questions, but as providing means to systematize questions and change perspectives among questions, which in the end can at least be powerful enough to diminish the importance of having answers.

3　Conclusions

The second edition of the workshop *Ethics and Morality in Business Informatics* was held at the International Conference on Wirtschaftsinformatik (WI) 2020 in Potsdam. As a whole, the contributions of this year's workshop have reinforced what the first edition of the workshop in 2019 has anticipated. The need to direct attention at ethical issues in the objects and methods of BIS research is today more palpable than ever. Among the reasons for this need are continuously rapid advances in IT and communication technologies, as well as growing disparities in the relative power of the various participants of a digitized economy. An interesting observation about this year's contributions is that many of them, in some way, apply traditional BIS techniques to the very methods of BIS research itself, seeking to heighten awareness of ethical issues bound up with the application and the products of these methods. We are looking forward to see how these efforts will be continued, and we hope to have shown that we as BIS researchers have both the responsibility and the ability to study the ethical implications of the IT artifacts that constitute the subject of our discipline.

Many colleagues have contributed to the success of this second edition of the Workshop on Ethics and Morality in Business Informatics. We wish to express our gratitude to the program committee members, who have carefully reviewed the received papers. We sincerely thank all authors who have submitted a paper; and we thank all presenters and participants of our lively workshop in Potsdam. Finally, we are grateful for the unfailing support of the organizing committee members of the WI 2020 conference.

References

1. Kant, I. (1996) Critique of Pure Reason. Unified Edition. Translated by Werner S. Pluhar. Hackett, Indianapolis, Cambridge. First published in 1781 and 1787.
2. Schutt, R. & O'Neil, C. (2013) Doing data science: Straight talk from the frontline. O'Reilly Media.
3. Ananny, M., & Crawford, K. (2018) Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. New Media & Society, 20(3), 973–989.