

11-4-2021

From PAPA to PAPAS and Beyond: Dealing with Ethics in Big Data, AI and other Emerging Technologies

Bernd Carsten Stahl

De Montfort University, bstahl@dmu.ac.uk

Follow this and additional works at: <https://aisel.aisnet.org/cais>

Recommended Citation

Stahl, B. C. (2021). From PAPA to PAPAS and Beyond: Dealing with Ethics in Big Data, AI and other Emerging Technologies. *Communications of the Association for Information Systems*, 49, pp-pp. <https://doi.org/10.17705/1CAIS.04920>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in *Communications of the Association for Information Systems* by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.



From PAPA to PAPAS and Beyond: Dealing with Ethics in Big Data, AI and other Emerging Technologies

Bernd Carsten Stahl

Centre for Computing and Social Responsibility

De Montfort University, Leicester, UK

bstahl@dmu.ac.uk

Abstract:

Researchers have long referred to privacy, accuracy, property, and accessibility (PAPA) framework in discussing ethical issues in information systems. While all four constituent components remain relevant, technical progress and technological integration in organizations and society in the intervening almost 40 years call for researchers to consider the acronym. In response to Richardson, Petter, and Carter's (2021) proposal to add the term "society", I suggest that extending the acronym in more than one dimension would be useful. For example, one could extend the acronym to include stakeholders (e.g., individuals, organizations or society) and system-use stage (e.g., input, processing, and output). The third dimension is the ethical issue, which still includes PAPA but can be supplemented with others, such as bias, power distribution, and others. Therefore, I suggest that we not only need to extend PAPA to PAPAS but that we need to go beyond a list of ethical issues to capture the richness and complexity in which ethics and information systems interact.

Keywords: Ethics, Information Systems, Big data, Artificial Intelligence.

This manuscript underwent editorial review. It was received 12/08/2019 and was with the authors for 5 months for one revision. John Venable served as Associate Editor.

1 Introduction

German speakers widely use “Papa”, which probably equates to the English “dad” or “daddy”, as an affectionate term for father. Thus, as a native German speaker, I was fascinated when I first came across Mason’s (1986) seminal paper in which he discusses information systems (IS) ethics while referring to the privacy, accuracy, property, accessibility (PAPA) framework. In the intervening years, I have had children of my own who sometimes call me “Papa”, which has strengthened my emotional attachment to the term. Against this background, I read Richardson, Petter, and Carter’s (2021) paper in which they use and develop PAPA as a way to refresh the IS discipline’s approach to ethical issues (and particularly to ethical issues that arise from big data analytics (BDA)). To continue the association between the PAPA framework and family life a bit longer, for most people, their enculturation into social mores—into what is right and wrong, allowed or not, praised or punished—undoubtedly starts in their family environment. At its best, parental guidance can help children discover their moral boundaries in a clear, supportive, and loving environment, which helps them to not only follow rules but also question and understand their purpose, underlying values, and justification and develop a virtuous moral character that allows them to become well-adjusted members in society. In reality, not all family environments foster this enculturation. Paternalism is only one term that denotes a problematic approach to ethics with roots in a particular type of family constellation. Thus, PAPA may have a positive influence on ethics—but we have no guarantee that it will. For example, maybe we also require MAMA, the female equivalent or maybe the family metaphor has too many limitations.

In their paper, Richardson et al. (2021) invite scholars to more broadly discuss IS ethics, an invitation I happily accept. I believe that their paper constitutes a timely and relevant reminder about the importance of ethics in the IS discipline. While IS researchers have a long history of discussing ethical issues (which Mason’s (1986) paper exemplifies), they have often treated ethics as a fringe topic. Of course, while researchers have made very visible interventions in the debate (Conger, Loch, & Helft, 1995; Mingers & Walsham, 2010; Smith & Hasnas, 1999), the debate in the IS discipline has typically focused on how organizations use digital technologies. In particular, questions such as which factors affect whether organizations adopt technology or which methods allow them to successfully outsource have been the most visible. One could argue that this situation is changing. Of course, how organizations use technologies remains an important topic, but digital technologies’ spread throughout all aspects of society means that nearly everyone feels their (sometimes positive but also sometimes ethically problematic) implications. Members of the IS research community, a community that has built up a rich repertoire of theories, methods, and insights into using computing technologies, find themselves in the privileged position of being able to contribute to reflecting on ethical issues and contributing to their resolution. Many expect some highly visible and widely discussed technologies, such as big data analytics, to have a disproportionate influence on humans’ wellbeing, which calls for a broader debate.

When talking about ethics, I find it worthwhile to discuss the term’s content. In English, it can refer to various related but clearly separate phenomena (Stahl, 2012). First, it may refer to socially agreed judgments about what people consider good or bad. Second, it may refer to justifications for why people would consider something good or bad, right or wrong. Third, it may refer to ethics as a philosophical discipline that provides theoretical accounts that can support judgments of right and wrong. In this latter category, we find the traditional ethical positions such as that an agent’s intention determines an action’s moral quality (as in deontology) (Kant, 1788), that an action’s moral quality depends on an action’s outcomes and consequences (as in consequentialism and utilitarianism) (Mill, 2002), or that an agent’s virtuous character constitutes the defining factor (as in virtue ethics) (Aristotle, 2007; MacIntyre, 2007; Vallor, 2016). Of course, other approaches exist as well, such as the ones that researchers have applied to moral concerns in the IS discipline (Bryant, Land, & King, 2009; Culnan & Williams, 2009; Elbanna & Newman, 2013; Martinsons & Ma, 2009; Walsham, 1996).

I briefly introduce ethics as a concept to highlight its complexity. To speak about “ethics” typically requires that one explain which aspect among the three I present above that one refers to. In many cases, however, scholars view ethics as a concept as unproblematic. Mason’s (1986) paper exemplifies this point: he introduces four issues as ethical with little reference to what makes them ethical issues. In discussing property, he refers to Mill, Bentham, Adam Smith and Kant but only to illustrate the fact that social concerns caused people to create ethical theories and without referring back to how their work can inform why one perceives an ethical issue as such or how one could deal with conflicts or tradeoffs.

This discussion brings me back to Richardson et al.'s (2021) paper. By adopting the PAPA approach, they focus on one particular angle of the ethics discussion. They begin by noting that BDA raises existing issues that Mason originally described to the fore but raises new ones that go beyond them. They demonstrate why and describe privacy, accuracy, property and accessibility as issues that BDA can exacerbate. They then argue that we need to complete these issues and suggest adding "society" to progress from PAPA to PAPAS.

I agree with their approach in several ways. As I indicate above, I think that the IS discipline has a good ability to engage in the debate about information technologies and ethics. Scholars can best facilitate such a debate by drawing on existing and established figures of thought, and revisiting PAPA represents one way to do so.

However, I believe they could further develop aspects of their approach. First, they could further develop a stronger link to the current ethics and technology debate. Ethics in digital technologies has gained prominence in academic discussions, in policy, and in the media. Richardson et al. (2021) do not draw on this discourse in the depth it deserves. Many publications have addressed ethical issues associated with big data and analyzed in detail how understanding big data can help researchers integrate ethical concerns in big data research (Clarke, 2018; Saltz & Dewar, 2019; Zook et al., 2017). In addition, a fast-growing number of policy- and industry-oriented publications have addressed such ethical issues (Accenture, 2016; Cabinet Office, 2016; International Bioethics Committee, 2017). Second, they could engage in more detail with the discourse on ethics of AI. Not all AI approaches require or use big data, but most high-profile cases that exemplify ethical concerns about AI involve machine learning that organizations typically implement through neural networks and rely on large amounts of training data. Big data analytics, on the other side, does not have to but often does rely on AI techniques to understand data. Thus, it makes sense to refer to AI debate ethics when talking about BDA ethics. And again, numerous academic publications (Cath, 2018; Dignum, 2018; Floridi et al., 2018; Floridi & Cowsls, 2019; Jobin, Ienca, & Vayena, 2019; Winfield & Jirotko, 2018) and an ever-expanding number of policy and industry interventions cover the ethics of AI (Elsevier, 2018; Executive Office of the President, 2016; House of Lords, 2018; Stix, 2019).

By engaging with these broader discourses, one could possibly preempt a potential weakness in Mason's (1986) approach; namely, their focus on particular ethical issues and, thus, their neglect for others. In critiquing Mason's (1986) approach, Fairweather (2003) pointed out that reliance on a code or on a list of ethical issues can be counterproductive if it engenders the assumption that successfully dealing with them will mean one has successfully dealt with ethics. Here I turn to my reservations about extending PAPA to PAPAS. I have two concerns. First, Richardson et al. (2021) seem to suggest that, by adding "society" to PAPA, they address a shortcoming in it, which leads straight back to Fairweather's (2003) objection about incomplete codes. Why would we think that adding these new concerns now covers all, or at least most, likely issues. Second, I think that "society" simply represents a completely different category from the original PAPA issues. One can easily see that privacy, access, property, and accuracy raise ethical concerns in the sense that people can imagine ways in which they affect rights and responsibility. Society simply refers to the greater context in which these issues become relevant. It also remains unclear why Richardson, Petter and Carter (2021) focus exclusively on the macro-level of society rather than the micro-level of the individual or the meso-level of the organization. I fully understand that we now face societal problems that Mason did not need to consider, but, rather than adding the term at the end, I would propose a matrix structure such as the one I propose in Table 1.

Table 1. Integration of Social Levels into PAPA

	Privacy	Accuracy	Property	Accessibility
Micro level (individual)				
Meso level (organization)				
Macro level (society)				

One could then fill in the individual cells using examples that illustrate how the various issues manifest at the different levels, which could provide links to possible ways to address them. While doing so would solve my second concern, it still looks closed and comprehensive, which is in itself problematic. In order to address this problem, I propose showing that the table has an open structure such as what I show in Table 2.

Table 2. Extension of Levels and Issues

	Privacy	Accuracy	Property	Accessibility	Bias	Power	...
Micro level (individual)							
Meso level (organization)							
Macro level (society)							
Other levels							

The two example issues I add in this table (i.e., bias and power) represent high-profile issues that likely relate to big data analytics. Of course, many more potential issues exist. These issues may also have a bearing on other levels beyond the ones I list, such as international organizations, semi-organized groups, and so on.

Richardson et al. (2021) discuss the issues by looking at the inputs, processing, and outputs, which I think is a good idea that can give clarity with regards to ethical issues' nature and possible ways to address them. Integrating this idea into the categorization suggested above would lead to a three-dimensional (3D) matrix such as the one I present below in Figure 1.

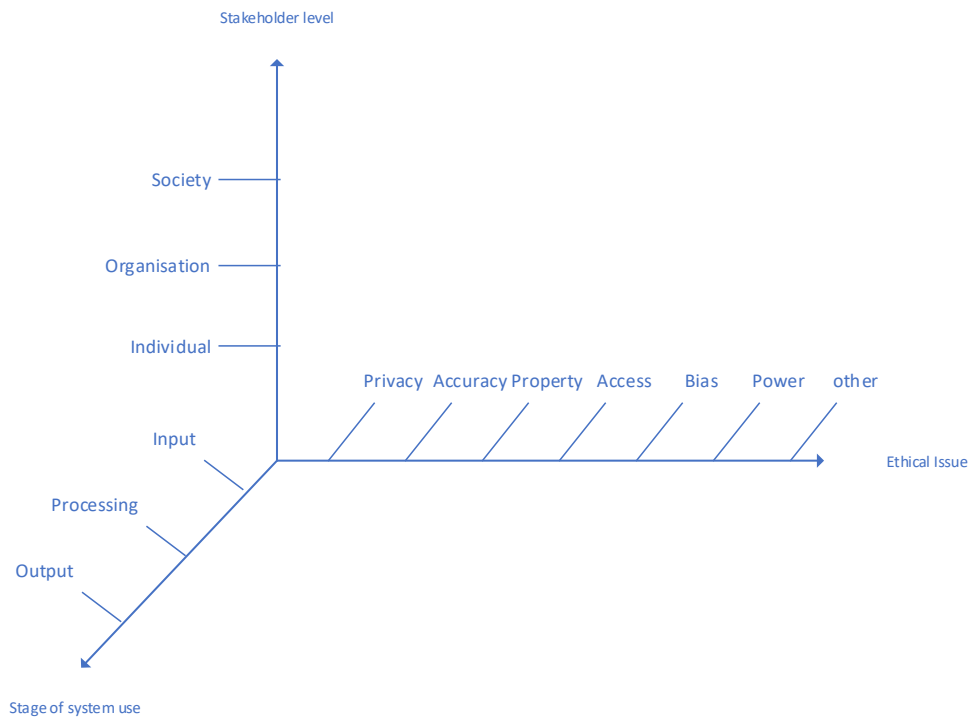


Figure 1. Space of Ethical Issues in IS that Comprises Three Dimensions: Ethical Issue, Stage of System Use, Stakeholder Level

Extending the PAPA model in this way has several advantages over the way that Richardson et al. (2021) propose. First, the model does not look closed and does not falsely suggest to IS scholars that they have addressed all ethical issues by working through PAPAS. Second, one can easily amend it to include additional aspects that related discussions, such as those around ethics and AI, have already explored. Third, it still builds on well-established ideas in the IS discipline and, therefore, will resonate with IS scholars.

If we accepted this approach to remodel the PAPA framework, we would next need to consider what we should do with it. In the first instance, I think we could use it to describe the ethical issues that exist. By using the 3D matrix, one could conduct a very fine-grained analysis of ethical issues. One could then use the understanding they obtain from such an analysis as the basis for a deeper ethical analysis, which brings us back to the ethics concept and the role of ethical theory. What makes an issue an ethical one, and how could one describe it from different theoretical perspectives? One could then use this more detailed understanding to map existing solutions. For instance, the current focus on ethics in AI seems to

have led to a large number of activities that try to establish ways to mitigate ethical issues while broadly ignoring that researchers in the IS and other disciplines have discussed it for decades. An ability to map ethical issues to existing mitigation strategies that range from legislation and regulation to standardization and professionalism all the way to development guidelines and personal virtue would be helpful in highlighting which issues most require attention. Furthermore, such insight would prove most helpful in determining a research strategy or roadmap that the IS community could discuss and that could be a call to arms for IS scholars to give ethics the attention the topic deserves.

The IS discipline tends to pride itself in its ability to understand practical problems and produce solutions that one can implement in practice. I hope that Richardson et al.'s (2021) work, maybe supplemented with the suggestions I make here, can help guide the discipline to deal with some of the biggest issues the world currently faces. PAPA alone may not be sufficient, but maybe we can build a broader family of ethical concerns and ways to address them to make sure that information systems' benefits justify their economic, social, and ethical costs.

References

- Accenture. (2016). *Building digital trust: The role of data ethics in the digital age*. Retrieved from <https://apo.org.au/node/71946>
- Aristotle. (2007). *The Nicomachean ethics*. Minneapolis, MIN: Filiquarian Publishing.
- Bryant, A., Land, F., & King, J. L. (2009). Editor's introduction to the special issue on ethical issues in IS research. *Journal of the Association for Information Systems*, 10(11), 782-786.
- Cabinet Office. (2016). *Data science ethical framework*. Retrieved from <https://www.gov.uk/government/publications/data-science-ethical-framework>
- Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society B*, 376(2133).
- Clarke, R. (2018). Guidelines for the responsible application of data analytics. *Computer Law & Security Review*, 34(3), 467-476.
- Conger, S., Loch, K. D., & Helft, B. L. (1995). Ethics and information technology use: A factor analysis of attitudes to computer use. *Information Systems Journal*, 5(3), 161-183.
- Culnan, M., & Williams, C. (2009). How ethics can enhance organizational privacy: Lessons from the ChoicePoint and TJX data breaches. *Management Information Systems Quarterly*, 33(4), 673-687.
- Dignum, V. (2018). Ethics in artificial intelligence: Introduction to the special issue. *Ethics and Information Technology*, 20, 1-3.
- Elbanna, A., & Newman, M. (2013). The rise and decline of the ETHICS methodology of systems implementation: lessons for IS research. *Journal of Information Technology*, 28(2), 124-136.
- Elsevier. (2018). *Artificial intelligence: How knowledge is created, transferred, and used*. Retrieved from <https://www.elsevier.com/?a=827872>
- Executive Office of the President. (2016). *Preparing for the future of artificial intelligence*. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf
- Fairweather, N. B. (2003). No, PAPA: Why incomplete codes of ethics are worse than none at all. In T. W. Bynum & S. Rogerson (Eds.), *Computer ethics and professional responsibility: Introductory text and readings* (pp. 142-156). New York, NY: Wiley.
- Floridi, L., & Cows, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*. Retrieved from <https://hdr.mitpress.mit.edu/pub/l0jsh9d1/release/7>
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707
- House of Lords. (2018). *AI in the UK: Ready, willing and able?* Retrieved from <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389-399.
- Kant, I. (1788). *Kritik der praktischen Vernunft*. Ditzingen, Germany: Reclam.
- MacIntyre, A. C. (2007). *After virtue: A study in moral theory*. Notre Dame, IN: University of Notre Dame Press.
- Martinsons, M., & Ma, D. (2009). Sub-cultural differences in information ethics across China: Focus on Chinese management generation gaps. *Journal of the Association for Information Systems*, 10(11), 816-833.
- Mason, R. O. (1986). Four ethical issues of the information age. *MIS Quarterly*, 10(1), 5-12.
- Mill, J. S. (2002). *Utilitarianism* (2nd ed., G. Sher, Ed.). New York, NY: Hackett.

- Mingers, J., & Walsham, G. (2010). Towards ethical information systems: The contribution of discourse ethics. *MIS Quarterly*, 34(4), 833-854.
- Richardson, S., Petter, S., & Carter, M. (2021). Five ethical issues in the big data analytics age. *Communications of AIS*, 49, 4-21.
- Saltz, J. S., & Dewar, N. (2019). Data science ethical considerations: A systematic literature review and proposed project framework. *Ethics and Information Technology*, 21, 197-208.
- Smith, H. J. & Hasnas, J. (1999). Ethics and information systems: the corporate domain. *MIS Quarterly* 23(1), 109-127.
- Stahl, B. (2012). Morality, ethics, and reflection: A categorization of normative IS research. *Journal of the Association for Information Systems*, 13(8), 636-656.
- Stix, C. (2019). *A survey of the European Union's artificial intelligence ecosystem*. Retrieved from <https://ec.europa.eu/jrc/communities/en/node/1286/document/survey-european-union%E2%80%99s-artificial-intelligence-ecosystem>
- International Bioethics Committee. (2017). *Report of the IBC on big data and health*. Retrieved from shorturl.at/bjJ36
- Vallor, S. (2016). *Technology and the virtues: A philosophical guide to a future worth wanting*. Oxford, MA: Oxford University Press.
- Walsham, G. (1996). Ethical theory, codes of ethics and IS practice. *Information Systems Journal*, 6(1), 69-81.
- Winfield, A. F., & Jirotko, M. (2018). Ethical governance is essential to building trust in robotics and AI systems. *Philosophical Transactions A: Mathematical, Physical and Engineering Sciences*, 376(2133).
- Zook, M., Barocas, S., Boyd, D., Crawford, K., Keller, E., Gangadharan, S. P., Goodman, A., Hollander, R., Koenig, B.A., Metcalf, J., Narayanan, A., Nelson, A., & Pasquale, F. (2017). Ten simple rules for responsible big data research. *PLoS Computational Biology*, 13(3), e1005399.

About the Authors

Bernd Carsten Stahl is Professor of Critical Research in Technology and Director of the Centre for Computing and Social Responsibility at De Montfort University, Leicester, UK (www.ccsr.uk). His interests cover philosophical issues arising from the intersections of business, technology, and information. This includes ethical questions of current and emerging of ICTs, critical approaches to information systems and issues related to responsible research and innovation.

Copyright © 2021 by the Association for Information Systems. 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 full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints are via e-mail from publications@aisnet.org.

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.