



HAL
open science

Ethics and AI Issues: Old Container with New Wine?

Fred Niederman, Elizabeth White Baker

► **To cite this version:**

Fred Niederman, Elizabeth White Baker. Ethics and AI Issues: Old Container with New Wine?. 20th Conference on e-Business, e-Services and e-Society (I3E), Sep 2021, Galway, Ireland. pp.161-172, 10.1007/978-3-030-85447-8_15 . hal-03648124

HAL Id: hal-03648124

<https://inria.hal.science/hal-03648124>

Submitted on 21 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Ethics and AI Issues: Old Container with New Wine?

Fred Niederman¹ and Elizabeth White Baker²

Saint Louis University, St. Louis, Missouri 63146, USA

² Virginia Commonwealth University, Richmond, Virginia, USA
fred.niederman@slu.edu

Abstract. This paper reflects on what differentiates AI ethics issues from more general concerns raised by all IS applications. Examination of the PAPA framework advanced in 1986 by Richard Mason suggests that the categories of problems remains much as they have over the decades, with the exception of a new set of societal issues. Within these categories, however, the increase in capabilities of computing generally and AI in particular, shift some affordances from only possible to realized and the ethical issues attached to these affordances come to the fore. AI, however, is a catalyst for deeper philosophical considerations about the nature of mind, thought, agency, and responsibility.

Keywords: ethics, artificial intelligence, AI, PAPA framework, Internet of behavior, distributed cloud, hyper-automation

1 Introduction

“Nowhere is the potential threat to human dignity so severe as it is in the age of information technology, especially in the field of artificial intelligence (p. 9),” Richard O. Mason (1986).

“Be nice to your vacuum cleaner, it may become more sentient than you before you have a chance to apologize for your decades of disrespect,” Fred Niederman (2021).

Literally thousands of scholarly articles have been written in the past decade alone about ethics and artificial intelligence (AI). Yet it is not clear how much has been added to an initial view of IS ethics produced by Mason (1986). This paper aims to review in broad brush what Mason’s framework reveals and what remains for further examination.

An ABI-Informs search on April 9, 2021, on the key words “artificial intelligence” and ethics yielded 27,821 retrieved articles. Narrowing down to scholarly journals yielded 3,152 articles. *AI Magazine* led the pack with 118 articles, *Communications of AIS* is among the leaders within the MIS community with 61 articles and the *Journal of AIS* led the “Basket of 8” designated by the AIS senior scholars’ college with 27. Special

issues have been devoted to the topic including ones in the journals: *Ethics and Information Technology*, *Philosophy & Technology*, *Transactions on Human-Computer Interaction*, *Proceedings of IEEE*, and *Journal of AIS* (Aggarwal, 2020; Benbya, et al., 2021; Dignum, 2018; Winfield, et al., 2019; Robert, et al., 2020). Interest in the topic extends beyond the discipline of IS to include scholars in the humanities, social sciences, computer science, and engineering.

Numerous books, in particular *Weapons of Math Destruction*, (O’Neil, 2016), have addressed ethics and AI, including sophisticated analytics. O’Neil’s book is particularly effective at pointing to impacts from AI and analytics and just how much these can vary between stakeholders -- providing significant convenience to a large number of consumers, while providing extreme harm to relatively few for whom the systems do not operate as intended. *Future Politics* by Susskind (2018) speculates at length about the potential effects of AI (among other technology influences) on relationships between individuals and institutions in the future. It proposes that society is likely to face emergent issues as a result of advancing technology, create new rules and institutions to respond to these, and suggests that such a future result in nightmare or enlightened scenarios for the people of those days.

Much writing about AI ethics pertains to relatively broad threats such as the erosion of a need for human labor characterized in some quarters as the ‘robo-apocalypse’, and the pernicious effect of inappropriate exclusion from modern society based on biased or erroneous information. Writings on these topics tends toward showing the dangers of such threats (Clarke, 2019) or proposing that they are unlikely to manifest (Willcocks, 2020), noting potential ways to avoid or ameliorate them. Others focus less on particular AI applications or dangers and more on preventing general ill effects through attempts to provide rules or codes of ethics on one hand or algorithms for preventing, detecting, and correcting errors. Siau and Wang (2020), in this regard, usefully differentiates AI ethics in terms of those applying to creating new AI versus the effects of AI on stakeholders.

In this essay, we consider whether and, if so how, AI presents new ethical challenges beyond those that have confronted information systems and reactions to them from AI’s beginnings. In this paper we address questions arising from reconsideration of Mason’s (1986) PAPA model (privacy, accuracy, property, accessibility). The overarching research question of this study is: What, if anything, differentiates AI-related ethical issues from all other information systems ethical issues? Secondary questions consider whether Mason’s framework is sufficient for organizing discussions about AI and ethics. If it is not sufficient, what remains to be added?

This line of questioning is based on the premise that if we have accumulated knowledge of ethics pertaining to information systems generally, some may be directly applied to issues relative to AI. Once applied, are there remaining issues to be addressed? Mason’s framework is certainly not a representation of the accumulation of all knowledge about IS and ethics, but it is a highly cited and fundamental piece worth

considering as an initial entry point to comparison of AI and general IS relative to ethics.

Underlying this discussion is the view that AI can be understood as a collection of techniques that to some degree seek to mimic human capabilities (e.g., language recognition, movement through robotics) or develop alternatives that substitute for human decisions and/or actions. The dividing line between AI and other information systems is fuzzy because; (1) AI components may be embedded in more traditional systems. For example, some AI programming may search for fraudulent transactions as a component of an enterprise or accounting system; (2) to some degree AI and analytics represent porous categories where some particular techniques could be claimed by either. For example, Bayesian statistics may be regarded as underpinning either analytic or AI techniques; and (3) historically many information systems applications, from COBOL implemented transaction systems to proprietary enterprise systems have substituted for human decisions and actions without concerns for replicating the manner in which humans performed these tasks.

Mikalf and Gupta (2021) methodically develop a list of capabilities pertaining to organizations relative to their ability to create applications using AI and, as a result of their use, create value. These capabilities are framed in terms of technology, human resources, and other organizational factors. They do not explicate AI capabilities in terms of their affordances or how they are put to use in the world. Thus AI is viewed in terms of organizational skill and technical capabilities rather than affordances such as: extrapolating past activity into probabilities for future performance, connecting ‘sensory’ data such as from vision or language to actions contingent on the content of such data, helping to eliminate infeasible choice options, or finding unexpected connections (say selecting team members from hundreds of thousands in a multinational enterprise for a special project) where limits of human experience and search time would make such activities impossible.

In the following sections, we review Mason’s framework, discuss what it does and does not do for advancing ethical considerations, consider particular issues that AI raises in the philosophical realm, particularly regarding ‘consciousness’; present three examples of emerging technologies and how the PAPA framework can identify issues but also leaves some remaining unclassified, which we propose falls under an additional ‘societal’ category, then we conclude with discussion about what remains for dealing with AI ethics beyond the framework.

2 PAPA Issues Then and Now

Mason (1986) produced a seminal categorization of ethical issues particular to IS. He called them the “PAPA” issues: privacy, accuracy, property, and accessibility. For each he provides examples. There is also a modest sequential nature to the issues, particularly as privacy leads into accuracy concerns.

2.1 Privacy

Privacy in the framework represents a combination of issues. The first pertains to the ability of individuals to choose what to conceal or reveal from selected or all others information about themselves. This may be for reasons of economic consequence such as not revealing a medical condition that would change costs of insurance (or even eligibility). It may alternatively be for pure personal preference such as revealing that one is addicted to reading Marvel comics (a badge of honor in some circles, a sign of immaturity to others). The second pertains to the ability of information regarding an individual revealed for one purpose to be applied to non-disclosed subsequent purposes. Mason is aware of the cumulative effect of multiple revelations which he summarizes as: “Each additional weaving together of my attributes reveals more and more about me (p. 6).” Importantly, his view was expressed in 1986 before such data integration has largely become a widespread fact of daily life.

In contemporary terms, we see not only the convergence of increasing numbers of data sets pertaining to individual behaviors and choices, but the increasingly sophisticated ability to combine these to extrapolate likely additional characteristics and to predict future choices and preferences. As a result, new affordances from chatbots to recommender systems shift from being possible to actualized.

Note that these issues of privacy do not assume any mistaken or malicious use of data. Rather even with non-distressing inputs, outputs may still cause harm when data has been combined. Issues where data is actually false or misinterpreted lead to the second element in PAPA, lack of accuracy.

2.2 Accuracy

Accuracy, or more properly inaccuracy, becomes a problem when (1) historic methods of interaction are subsumed by digital ones such that historic archives may not be recognized in the digital world – if there is no physical backup and the computer says you are a deadbeat, you have become one whether you have paid your obligations or not. When the data becomes the truth, inaccuracy can produce distressing results when translated back to the ‘real world’; and (2) people acting on data that should be but is not correct result in costly mistakes. Mason illustrates with faulty GPS information leading to a transportation accident with attendant costs. Mason notes the growth in such reliance: “Today we are producing so much information about so many people and their activities that our exposure to problems of inaccuracy is enormous (p. 8).”

Given that it is nearly impossible and relatively expensive to maintain perfectly clean data sets, much attention on accuracy pertains to finding and correcting those inaccuracies that make a difference. Note the attention in data warehousing to the cleansing of data being moved from independent transaction to integrated systems. In some scenarios all individual pieces of data may be correct (or at least not verifiably wrong), yet in combination lead to a poor conclusion. This can happen for example if a flawed procedure, say a contaminated blood sample, produced multiple readings all

of which are incorrect, but correctly recorded relative to the reported test results. Even accurate data may suggest membership in a group where such combinations are the norm, but where an individual does not conform. Hofstede (1991) warns very severely about the ‘ecological fallacy’ of generalizing from the group to each individual. Cases of the harm that can arise from such invisible sources are dramatically documented by O’Neil (2016).

2.3 Property

Property pertains to the ownership of information as well as the tools that produce and manipulate it. Generally intellectual property refers to knowledge embedded in artifacts, but the content and data itself should be included as well. Mason focuses, among the many things which content ownership implies. But consider one producing a political tract being paired by an information consolidator, with an opposing tract which contains misinformation and personal attacks that one doesn’t have the opportunity to refute. Even if not a comma is changed, the meaning and purpose of the initial tract may be significantly misused.

The systems we interact with now tend to be combinations of pieces supplied by an array of vendors which operate together to produce particular outcomes. However, when there are snags, it is often difficult to match the source of the problem to one particular component – particularly if each is operating as expected but are out of mutual alignment. We might call this the responsibility for floating and ghost glitches.

2.4 Accessibility

Accessibility refers to the capacity of individuals or organizations to acquire and use data. Mason posits the keys to this are three components of computer literacy: the intellectual skills to handle information; access to the technologies that connect to that information, and access to the information itself. It is noteworthy that the first two of these require economic and psychological investment by those who would seek access and the last pertains to the legitimate and illegitimate efforts to shield information from being acquired freely. Mason focuses on inequalities in the ability to pay these prices: “Thus the educational and economic ante is really quite high for playing the modern information game. Many people cannot or choose not to pay it and hence are excluded from participating fully in our society (p. 11).”

Over the years this situation has come to be called the ‘digital divide’ which marks differential access by economic status; location or ethnicity of origin; rural or urban; gender; age; physical capabilities and various other characteristics. In a world strongly moved, if not dominated, by economic performance, what would motivate vendors of software or hardware to produce versions of their products and services for subsets of customers where revenues generated are unlikely to payback the costs of modification for special needs?

3.0 PAPA Implications for AI Ethics

Before turning to the relationship of AI ethics issues and PAPA, it is worth considering what the framework is and what it isn't. As a framework, PAPA is a categorization scheme, or taxonomy, if one prefers. Following Bailey (1994) there are generally two types of taxonomy – inductive and deductive. Inductive taxonomies sort instances into groups where each group represents a category. Quantitatively, such a procedure may begin with a set of instances and use mathematical techniques like cluster analysis to sort the instances into types. It becomes a challenge to scholars to both label and ascertain the key differentiators among categories.

In contrast, the deductive approach starts with logically derived categories, sometimes by observing variations on dimensions thought to be meaningful such as aggressive/passive; communicative/quiet and using placement on both measures to segregate instances into categories. An advantage of this approach includes straightforward assignment of instances into categories. On the other hand, it does not necessarily take into account differences on other dimensions that may be equally or even more relevant, variance among instances within categories, and dimensions where evaluation of an instance is not clear.

3.1 Uses for PAPA

Once a taxonomy is established it can have multiple uses. One purpose is to create a basis for classifying instances. Given a particular instance, like drones falling out of the sky and injuring pedestrians, we might use the PAPA framework to consider what sort of issue this is. In this case, I'd surmise it would be likely to be considered a "property" issue in terms of who owns the drone and who owns liability for its proper usage. The classification relative to the framework in itself would be trivial (who cares what sort of problem it is when I need someone to drive me to the hospital?)

However, if we can assign an instance to a category, we can envision that the instance may inherit the general characteristics of that category. If the falling drone is a property problem it may inherit issues relative to varied stakeholder contributions that are present in many cases such as designer, builder, operation, and owner each having some potential responsibility. To the extent that such a premise is the case for all property issues, identifying the falling drone as a property issue highlights at least a first set of considerations. In this way categorization can create value by shortcutting the initial analysis process when a new instance is identified such that the general attributes become known and more specific characteristics of the instance can become the focus of further investigation. Note, however, that such initiation of discussion cannot substitute for the independent analysis of each emergent case. The drone falling from the sky may have some analogy to a self-driving car crashing into a pedestrian, however it also has differences including the potential for a different kind of damage and different level of victim awareness (e.g., as a pedestrian, one knows, or should know, that the space is shared with automobiles driven by humans or not; but the sunbather cannot be expected

to anticipate a drone falling from the sky – at least at this time when drone use remains rare).

Although Mason (1986) discusses each of the categories in broad terms and presents examples, it is not clear whether any standard properties have been defined that can be applied to surfaced instances within the category.

Another use of a categorization scheme like PAPA is to help anticipate emergent issues. As technologies grow near creating new affordances, we may use the PAPA framework to anticipate classes of emergent issues. We know that there are property issues relative to human driven autos, but as we approach driverless ones, we can see that the list of contributors to the new technology grows to include programmers and other IT specialists along with the manufacturer and driver. Adding GPS to the auto (whether driven by humans or not) adds another dimension of affordances – like automated routing for the driver and recording of locations for archivists or whomever else can and wants to wrangle access) suggests a range of privacy, accuracy, and access issues.

As new technologies and platforms emerge, from fintech to crowdsourcing to gamification and well beyond, the PAPA framework can be applied to anticipate emergent general ethical instances such that designers and builders may be able to design them out of systems and/or processes for detection and remediation may be set in advance. Further, periodic consideration of PAPA categories can be used relative to the interaction and accumulation of computing technologies to consider emergent ethical instances that do not adhere to one and only one emergent technology. For example, a growing sense of anxiety may derive not just from social media, ubiquitous computing, and the move toward digitization of many previously physical processes but from their accumulation. Each may provide benefits over cost worth their adoption but collectively they may change society in ways that generate new concerns.

Note that this discussion does not address AI as separate from other computing technologies. AI however is worthy of consideration as a generator of new applications or affordances that create new ethical instances. It is also an accelerator of new affordances in combination with other technologies. AI is a component enabling technologies like drones and self-driving cars, but both problems and solutions are intertwined with other systems (e.g., better sensors, coding for collaboration whether using AI algorithms or not, and even physical systems like lighter materials and better brakes).

3.2 What PAPA Doesn't Do

The strength of PAPA is in helping us understand and plan. That is not to say that understanding say a privacy issue with Alexa, or the other home companion devices,

can positively lead to solutions through designing the problem out or providing countervailing tools for users. On the other hand, the PAPA framework still has some significant limitations.

PAPA does not offer any hints about how to deal with tradeoffs among stakeholders. For example, one might argue that I have a right to record every word you say to me, but equally you might have the right to speak without your words being recorded. Your right to the privacy of your words competes with my right to hold and review all that is presented to me. I suggest that nearly all issues that pertain to multiple stakeholders and present complex balances among competing reasonable if not valid interests. Knowing the category of this class of ethical issues may help us select a process for negotiating a compromise solution, but handling of such tradeoffs remains.

PAPA does not reduce uncertainty. A consequentialist ethics evaluates actions based on their results. Sometimes this might be quantifiable as a simple cost-benefit analysis, other times intangibles do not allow this. At the time of decision and action we can only estimate or guess what the outcomes will be. In the early 2000s when wireless systems were first being installed in coffee shops, there was widespread concern about pirating of the data of users of these public systems. To my understanding such concerns have not manifest (perhaps because they were highly publicized and of sufficient concern) in part because of technologies such as frequency changing but perhaps also because there are so many other easier, less costly, and less risky ways to steal data. Where the PAPA framework helps to anticipate issues and where these are viewed as sufficiently concerning, the process whether organized or ad hoc may indeed influence decisions and actions in a way that averts negative incidents.

PAPA does not resolve differences in moral values. Even if we could know the outcomes of any particular decision and action in advance, it is unlikely that agreement will be universal about which results are “good” versus “bad”. Algorithms that discriminate by postal code in assigning eligibility or cost of home loans may cause immense difficulty for those living in some neighborhoods. Some will view this as atrocious malfeasance, others as simply reflecting the underlying discrimination of society such that these differences are simply passed on to the AI system and, therefore, not its responsibility. PAPA can surface this sort of issue, but not resolve it.

3.3 AI and Consciousness

Another type of ethical consideration that PAPA does not directly address pertains to the issues emergent relative to the hypothetical consciousness and agency potentially embodied in machines. This is largely a philosophical question with varied positions espoused, often with quite a bit of certainty. On one hand is Dreyfus (1992) who maintains, with much more thorough discussion that consciousness is and must be embodied. On the other hand, Rorty (2018) and like-minded philosophers argue that what we experience as consciousness is an emergent property of physical brain processes and

therefore with no inherent reason why other sources of elemental processing should not also create consciousness as an emergent property.

Let us suppose, for example, that we have come to a point where creating machine consciousness can be achieved. Should it be done? From a universal perspective, the most important principle is uncertainty. We cannot begin to predict the results even if we can anticipate some and are fairly certain that others will emerge. If we take this step, even with a backdoor to ‘pull the plug’ if things go wrong, could we be justified ending the life of another sentient being, even if sentient by virtue of our own choices and actions? If the capability to create consciousness is near, would there be any way to stop absolutely everyone capable from deciding individually whether or not to go forward?

Another issue pertains to the rights and responsibilities of sentient machines. In many, if not all, human societies we grant some rights to animals of various sorts, even if not equivalently to humans. We recognize to varying degrees the consciousness and contribution of animals from endangered species to work animals from seeing-eye dogs to horses that pull wagons carrying tourists through recreational centers. Assuming varied possible levels of intensity in extending human rights to machine agents, ranging from full human status to regulations on maintenance, we might meaningfully ask questions about the implications for human ethics. Following PAPA, how might knowledge obtained by sentient machines affect our privacy? Do we have the right to expunge information gathered before sentience occurred? Assuming dozens, hundreds, or millions of such machines it is almost guaranteed that they will have varied “experiences” and form a wide range of inferences about us that will vary from the affirming to the very negative. If the machines are built from code written by individuals relative to various license agreements drawn for example from open source libraries, do the authors of the code have claims on benefits created by the machines or liability for claims made against them?

On the other hand, if calls for strict ethical screening of the processes for building AI and the data sets on which they are trained are instituted, we may find ourselves as humans in a position where our tools are more trustworthy than our colleagues. Perhaps this has long been the case as, simplistically, I’d trust a jack to hold the car up while I am working on it than even three strong and well-intended friends. It is a little much to believe we can program effective ethical guardrails on our AI tools when we demonstrably have so much trouble both defining and abiding by even simple ethical guidelines among ourselves. If you doubt this, check the local newspaper for listings of murders, embezzlements, and domestic violence – events that would rarely be considered highly ethical even without the help of AI robots. God, Allah, and/or Atman please help the vacuum cleaner that can feel pain before it can defend itself.

We may make the assumption that such machine consciousness is far off in the future, but when is the right time to start considering these issues? Will we as humans

even know when machines develop sentience any more than we know how sentient are whales or mice?

4.0 IS PAPA sufficient?

Given the early stages of the development of a new technology, one might use the framework proactively to project possible ethical issues following the PAPA framework. But we should not take this for granted but rather test how well it does work. Let us consider three currently emergent technology trends as proposed in the 2021 Gartner strategic trends listing (Pannetta, 2020): internet of behaviors; distributed cloud, and hyper-automation (see Table 1).

Table 1. Three emerging technologies viewed through PAPA categories.

	Internet of Behaviors	Distributed Cloud	Hyper-automation
Privacy	Integration of disparate data sources Exposure of individual items and/or inferences	Forced privacy if data is not compatible across platforms	???
Accuracy	Tolerance of sensors to error Compounding of error through integration	Finding data where it is spread across physical media Keeping all data updated simultaneously throughout system	Inability to intervene mid-process to avoid or correct mistakes
Property	Locus of decision making for access to each thread as well as collection	Claims of host toward ownership of data Responsibility of host for effects of false and harmful data	Concentration of ownership in fewer hands
Access	Authority and ability to interpret data	Chokepoints and potential arbitrary host rules (and changes)	Exclusion of key stakeholders from usage
Societal Well Being	Movement toward “surveillance society”	Faster, more efficient distribution of data and content	Influence on number and type of jobs; the nature of work

It is relatively easy to map issues of internet of behaviors to the PAPA framework. Issues generated by the internet of behaviors follow and extend especially concerns for privacy, accuracy and access. These might be closely related to the “surveillance society” as described by Clarke (2019) where hyper-automation refers to the arrangement of work and, more broadly, to the accumulation and distribution of resources. It can be illustrated with the metaphor of adding or taking away chairs in the game ‘musical chairs’ or changing rules which favor one group (perhaps those with more STEM education) relative to others without the same capabilities. It is not clear the extent to which AI technologies and codes are required to enable the internet of behaviors. To the extent that such implementation is agnostic to particular coding approach, AI may have no strong relationship to ethics in this arena.

Relating distributed cloud and hyper-automation to the PAPA framework is more difficult, perhaps because they are of a more infrastructural technical nature, but not impossible. Distributed cloud is mostly about where data is stored, though it could have effects on users by enabling easier access and perhaps greater security from external threats by affording larger investment in preventative tactics. If having most relevant data nearby and access to the full range of data from across the entire “cloud” creates some new affordances, these might generate new ethical issues. If integration across separate clouds (e.g., data on Amazon is not accessible from Google, or vice versa), then perhaps issues involving the selection of host and either the forcing of ecosystem choice or expenses for operating on multiple platforms might actually constrain some new innovation while presenting their own difficult business choices. Issues of hyper-automation would concentrate on the basis for ownership of systems and access across stakeholders to determining its use.

This set of ethical issues pertains less to individuals (except perhaps in the numbers affected) and more in structural shifts in the community. For lack of a better term, we might label the category of these issues; societal well-being. In our view this term pertains to the changing environment within which individuals might operate. It can be visualized as problems that may arise when each component is optimized (or at least operating in a satisfactory way) and approved but where their accumulation changes the milieu in which we operate. For example, we may value surveillance of public spaces to reduce crime but if the surveillance becomes too pervasive, social life may become overly artificial and fail to accommodate its users.

4.0 CONCLUSION

Our moral imperative is clear. We must insure that information technology, and the information it handles, are used to enhance the dignity of mankind. To achieve these goals, we must formulate a new social contract, one that insures everyone the right to fulfill his or her own human potential (Mason, 1986, p. 11).

I'm afraid that like many discussions of AI and ethics, we raise more questions than we answer. That said, we think this brief essay suggests that the PAPA framework holds up reasonably well as a basis for analysis of AI ethical issues. It serves as a way to sort various possible negative effects of new AI generated affordances; it serves as a way to anticipate some emergent ethical issues as new technologies emerge. That said, it has some limitations. It focuses much more strongly on effects on individuals. In considering particular emergent technologies, we surface one category of ethical concern, social well being, that highlights a different set of ethical issues, or suggests another way to look at them.

Overall, we believe that the ethical issues faced by the emergence of AI are largely extensions of those generated by the continued growth of the technical capacities of information systems and, as a result, the growing storehouse of affordances. That said, the main effect of AI is to accelerate capacities and affordances such that entire families of issues move from hypothetical possibility to actualization.

In some quarters there is much faith that the solution to technical problems is more technology. Arthur (2009) seems to suggest this in documenting a view of the history of technology pointing out that many characteristics of sophisticated technology are the solutions to problems created by prior versions. To the extent that machine learning can reinforce underlying social discrimination, an algorithm to search for and remediate this would exemplify a technology solution to a technology created problem. But there is also pushback on this idea as exemplified by Kolbert (2021) describing the great environmental problems generated by engineering solutions to earlier problems. Algorithms that catch statistical discrimination may relieve the problem for individuals but may leave the underlying disadvantages causing the social issues, like discriminatory housing opportunities to remain unaddressed.

At the end of the day, we should be grateful to Mason (1986) for creating a useful framework, but not be satisfied with only categorizing issues. Now, more than 30 years following the presentation of this framework, it is difficult to say we've made positive progress in balancing competing values relative to the PAPA issues or developed ways to prevent or remediate harm when it occurs. On a personal note, having lived through a good portion of the information age since first diffusion of mainframe computers to the present (May 2021), it seems we are pretty good at technical solutions to technically created problems, we are good at adjusting to incremental changed expectations relative to privacy and evolving social norms, and at waiting out problems which eventually resolve themselves (although those harmed by them may not see it this way). On the other hand, we can use progress on methodically surfacing ethical issues emergent with new affordances, ways to quickly recognize and remediate harm done, and techniques for addressing varied stakeholder positions, overall uncertainty, and differences in values and moral opinions.

References

1. Aggarwal, N. Introduction to the Special Issue on Intercultural Digital Ethics. *Philosophy & Technology*; 33(4): 547-550. (2020).
2. Arthur, W.B. *The nature of technology: What it is and how it evolves*. Free Press. New York. (2009).
3. Bailey, K. D. *Typologies and Taxonomies: An Introduction to Classification Techniques*, (Vol. 102), Sage. (1994).
4. Benbya, H., Pachidi, S., and Jarvenpaa, S. Special Issue Editorial: Artificial Intelligence in Organizations: Implications for Information Systems Research. *Journal of the Association for Information Systems*; 22(2): 281-303. (2021). doi: 10.17705/1jais.00662
5. Clarke, R. Risks inherent in the digital surveillance economy: A research agenda. *Journal of Information Technology*. 34(1) 59–80. (2019). DOI: 10.1177/0268396218815559
6. Dignum, V. Ethics in artificial intelligence: introduction to the special issue. *Ethics and Information Technology*; 20(1): 1-3. . (2018)
7. Dreyfus, H. *What computers still can't do: A critique of artificial reason*. Cambridge, Massachusetts, MIT Press. (1992)
8. Hofstede, G. *Cultures and organizations: Software of the mind*. London: McGraw-Hill. (1991).
9. Kolbert, E. *Under a white sky: The nature of the future*. New York. Crown. (2021).
10. Mason, R.O. Four Ethical Issues of the Information Age. *MIS Quarterly*. Vol. 10 Issue 1, p5-12. (1986). DOI: 10.2307/248873.
11. Mikalef, P. and Gupta, M. Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management* 58(3). (2021): 103434.
12. O'Neil, C. *Weapons of Math Destruction*. Crown Publishing Group a division of Penguin Random House, New York. (2016).
13. Panetta, K. Gartner Top Strategic Technology Trends for 2021. <https://www.gartner.com/smarterwithgartner/gartner-top-strategic-technology-trends-for-2021/> Last accessed 2021/4/21.
14. Robert, L.P, Jr; Bansal, G.; Melville, N.; and Stafford, T. Introduction to the Special Issue on AI Fairness, Trust, and Ethics. *AIS Transactions on Human-Computer Interactions*; 12(4): 172-178. (2020).
15. Rorty, R. *Philosophy and the Mirror of Nature*. First Princeton Classics Edition, Princeton, New Jersey, Princeton University Press. (2018).
16. Siau, K. and Wang, W. Artificial Intelligence (AI) ethics: Ethics of AI and ethical AI. *Journal of Database Management*, 31(2), 74-87. (2020).
17. Susskind, J. *Future Politics*. Oxford University Press. Oxford, United Kingdom. (2018).
18. Willcocks, L. Robo-Apocalypse cancelled? Reframing the automation and future of work debate. *Journal of Information Technology*, 35(4) 286–302. (2020). HYPERLINK "" <https://doi.org/10.1177/0268396220925830>
19. Winfield, A., Michael, K., Pitt, J., and Evers, V. Machine Ethics: The Design and Governance of Ethical AI and Autonomous Systems, Special Issue Proceedings of IEEE. 107(3) 509 – 517. (2019).