POETs and Quants: Ethics Education for Data Scientists and Managers

M. Lynne Markus, Bentley University, mlmarkus@bentley.edu
Marco Marabelli, Bentley University, mmarabelli@bentley.edu
Xiaolin (Christina) Zhu, Renmin University, zhuxiaolin@ruc.edu.cn

Please cite as follows: Markus, M. L., Marabelli, M., Zhu, Xiaolin (2019). *POETs and Quants: Ethics Education for Data Scientists and Managers*, presented at the 3rd RICK workshop of Information and Organization, Cambridge, UK.

Our university, like many these days, is aggressively developing programs to prepare our students to meet the growing demand for expertise in data analytics and artificial intelligence. Aware of the many ethical concerns raised by data and algorithms (Newell and Marabelli 2015, Markus 2017), we began exploring the literature for ideas about what kinds of ethics education should accompany technical training in data analytic skills. This essay documents our reflections on the literature and our current thinking about analytics ethics education and the possible role of Information Systems scholars in developing and delivering this curriculum.

The first stage of our project was a broad and eclectic scan of the many literatures that might bear on our questions. A non-exclusive list of the topics we skimmed includes: philosophy, moral reasoning, computer ethics, information ethics, AI ethics, data science, business ethics, engineering ethics, ethics in information systems, biomedical ethics, ethics education, organizational misconduct, governance and regulatory compliance, corporate social responsibility, business law, and specific issues such as privacy, security, discrimination, predictive policing, and workplace monitoring. This highly non-systematic review inspired a number of conclusions, all of which are debatable and should be discussed in the Information Systems community. Key among those conclusions are two: 1) the need for and appropriate focus of ethics education are contested, and 2) current approaches to AI ethics education neglect two critical issues that the Information Systems community knows a great deal about—the educational needs of general managers and the dynamics of organizations. We conclude this essay with some preliminary thoughts about a curriculum on AI ethics aimed at general managers.

The need for and appropriate focus of ethics education are contested

Interest in ethics education waxes and wanes and always seems to spike after public scandals. Although business ethics has been a focus of research and teaching for decades, the topic experienced a surge of attention in the U.S. in the early 2000s after the misdeeds of Enron and WorldCom executives led to passage of Sarbanes-Oxley Act. Similarly, concerns about the ethical issues of big data use grew sharply in response to Snowden's revelations, several highly publicized data breaches, and Facebook's many violations of public trust.

In consequence of public outrage over scandals, ethics education is often viewed as the attempt to prevent individual misconduct by inculcating values, building character, and enhancing ***moral reasoning***. A typical educational strategy is to inform students about ethical theories such as utilitarianism (non-Western value systems are sometimes included) and to provide students with practice in applying these theories to particular situations.

Naturally, there is considerable skepticism about whether personal values and character can be taught at the university level. This skepticism fuels the inevitable curriculum fights that pit needs for technical skills against softer professional skills. It also favors the claims of those who would treat ethical concerns as a matter of regulation and compliance.

Of course, ethics education can be, and often is, viewed through lenses other than that of individual virtue and moral reasoning. Perhaps the major alternative approach is that of ***applied professional ethics***. Almost every occupational group that self-identifies as a profession (e.g., engineers, statisticians, computer scientists, and biomedical researchers) or that has experienced public opprobrium (e.g., internet marketeers, Wall Street bankers, credit bureaus, etc.) has developed ***codes of professional conduct*** and ***ethics educational curricula and materials***.

Most professional codes of conduct emphasize "virtues," that is, principles (e.g., honesty or fairness) that practitioners are expected to enact in their everyday work. There are many such codes, and there is even a small literature that compares, contrasts, and synthesizes ethical codes and that criticizes or defends them. In the AI ethics area alone, there are said to be as many as 80 published ethical statements by professional associations, governments, and NGOs (Jobin, Ienca et al. 2019). As one example, the ACM advocates the ethical principles of Fairness, Accountability, and Transparency for those who develop machine learning algorithms. While professional codes of conduct endorse high norms of practice, they are frequently short on practical guidance about how professionals should behave when values conflict or principles come under attack.

University-based ethics education often goes beyond codes of conduct to focus on specific practical problems that professionals encounter while plying their trade. For instance, engineers working on large construction projects are likely to experience pressure to accept or give bribes. Statisticians sometimes face demands from clients to alter their analyses or findings. Biomedical researchers frequently confront dilemmas around obtaining informed consent from patients. Business ethics curricula vary considerably, depending in part on whether they are taught by philosophers, business lawyers, or corporate social responsibility experts. Business ethics curricula commonly include attention to specific legal/ethical issues like deceptive advertising, customer privacy invasion, unlawful discrimination, and workplace surveillance. Pedagogical approaches in these curricula include readings that highlight contrary points of view on identified ethical issues, cases that bring ethical concerns to life, and scenarios for role playing and debating particular ethical decisions.

A strikingly different approach to the teaching of business ethics is Mary Gentile's Giving Voice to Values (Gentile 2010), which is premised on the idea that emerging business leaders readily recognize ethically questionable situations but lack skills in speaking up effectively about their concerns to appropriate organizational actors. Here, the curriculum also employs cases and role

playing, but the emphasis is not so much on the specific ethical issues in the cases, but on the soft skills needed to diagnose organizational politics and to confront people effectively.

In short, in our broad survey of different domains, we observed little consensus on whether, why, and how to teach ethics.

Current approaches to AI ethics education neglect the needs of general managers and the dynamics of organizations

When we looked in more detail at the literature on AI ethics, we drew two debatable inferences. The first is that the target audience for ethical guidance and education is largely specialist occupations such as data scientists, policy makers, and regulators. ***General managers*** in business and not-for-profit enterprises seem largely to have been ignored. The second was that the ethical guidance provided did not account for the ***organizational dynamics*** of developing and deploying algorithms and analytic insights.

Our reading of the AI ethics literature suggested to us the existence of two major author camps with different target audiences. One author camp consists of lawyers, ethicists, and social scientists who focus on digital technology's impacts on society and vulnerable social groups. These authors make policy recommendations aimed at government leaders, legislators, and big tech company executives. Examples of the advocated policy reforms include chain-link confidentiality in data sharing contracts, due process mechanisms for people adversely affected by automated decisions, and requirements to employ consumer opt-in instead of opt-out provisions for sharing customer data.

A second author camp consists primarily of computer and data scientists and specialists in computer ethics. These authors offer ethical codes and guidance to technical specialists, including academic researchers, and devise technological solutions intended to mitigate problems like security breaches, privacy invasions, and harmful discrimination. Examples of the kinds of solutions proposed by this author group include anonymization techniques, automated approaches to detect concept-drift in data streams, and blockchain-like strategies for tagging data elements with reuse restrictions.

Two things struck us as largely missing from these literatures. The first is the target audience of general managers of not-for-profit organizations (such as social service agencies and local and regional governments) and businesses in non-tech industries (such as retailing, financial services, sports, and health care). Many of these entities have already begun developing and deploying AI and analytic insights. Decision-makers in these enterprises need to be aware of the ethical issues involved in using data analytics and algorithms, because they will make key decisions about the features of the technology their organizations will build or buy and about how to deploy the technology and change business practices as a result of algorithmic inferences.

We note that the neglected target audience of general managers is one that Information Systems scholars have long addressed with educational offerings designed to improve their tech savvy. We have educated them, for example, in appropriate ways to approach technology selection, outsourcing, and project management decisions. As we IS scholars build more emphasis on

analytics and AI into our curricula, should we not also help general managers learn about the related ethical issues?

As we saw it, a second thing missing from the AI ethics literature is awareness of the organizational dynamics of technology development and deployment. Much of the literature seems to assume a social sphere consisting of two generic actor groups—"developers" and "users." Developers appear to be big tech companies and individual technical specialists, without regard to their possible employment in non-big tech enterprises or in universities. Who or what is considered to be a user is even more unclear. Is "the user" of a predictive policing algorithm a public safety enterprise, a lead user (who may in fact be the algorithm developer or maintainer), an information officer (an administrative assistant who interprets algorithm output into a simpler form for communication to police officers), or the police officers themselves? Would not the ethical issues faced depend on the "user's" organizational role?

That question highlights the organizational complexity involved in many deployments of analytics and algorithms. The organizational complexity of technology introduction is something that Information Systems scholars routinely study, because we know that the available actions differ at every step in the process, and that actions at any step can affect the outcomes experienced for better or worse. We conjecture that the absence of organizational dynamics from the literature on AI ethics is a regrettable and consequential omission (Markus 2016). Unless decision-makers understand the organizational context and politics of analytics development and deployment, they will not be able to take effective action to prevent analytic harms and to promote beneficial outcomes. Should not we IS scholars attempt to fill these educational gaps?

In short, we concluded that the AI ethics literature is missing two key foci that Information Systems scholars are well-positioned to address—an understanding of the needs of general managers and knowledge of organizational dynamics.

A proposed analytics ethics curriculum for general managers

As we see it, the most important things that general managers need to know about the ethical issues around data and analytics are these:

- Attending to ethical issues is sound business practice aimed at achieving business value as well as preventing bad business decisions, organizational misconduct, and reputational damage.
- Outcomes can be improved and some problems can be avoided by managers who understand technology development and deployment, practice a degree of foresight, and make appropriate decisions regarding the selection, design, and implementation of analytic technologies and redesigned business processes.

Put differently, general managers need to know that introducing analytic insights into organizational processes is a ***sociotechnical*** intervention, not purely a technical one. This means that managers must learn, not just about the ***technical*** issues, but also about the ***organizational*** and ***people*** issues that data analytics involves. This theoretical framework, sometimes known by the acronym TOP, pervades Information Systems textbooks today. With some updating and at least one key addition (***environment***), it can serve as a framework for the analytic ethics education of general managers as well.

Technical

General managers need to know, not just about the technologies of data analytics (e.g., Hadoop, R, data lakes, algorithm types), but also about:

- How use of these technologies can *promote negative outcomes* such as privacy invasion, unlawful discrimination, and stalking
- How use of various technological remedies (e.g., anonymization) can *reduce, but not eliminate*, negative outcomes from occurring
- The challenges of making algorithms *intelligible* and the needs of various stakeholder groups for effective explanations
- How important it is to monitor and *update* the performance of algorithms periodically given changes in data streams and human reactivity
- The technical and organizational *infrastructures* required for high quality and ethical uses of data and algorithms
- How to develop and assess a *use case* for applications of artificial intelligence and machine learning
- How to read an algorithm *data sheet* to compare alternative models and/or vendors

Organizational

General managers need to know, not just about the role of business process redesign in achieving good organizational outcomes from analytics, but also about:

- How the *organizational structure* of responsibilities for data science, information technology, regulatory compliance, etc., can *promote or inhibit* legal and ethical behavior
- How dealing with *external business partners* (for data, for algorithms, and for business services) can introduce additional challenges for regulatory compliance and ethical behavior
- How to develop and evaluate *business cases* for advanced automation
- How to employ *creative problem-solving techniques* to improve the business and ethical outcomes of proposed technical solutions

People

General managers need to know, not just about the possibility of "user" resistance to change, but also about:

- The tendency of technologists and managers to exhibit "*solutionism*" (inappropriate or excessive pro-technology bias such as the belief that algorithms are objective while humans are fatally flawed decision makers) and to be insufficiently aware of regulatory constraints and possible bias in their data and algorithms
- How *organizational culture* can promote or inhibit ethical behavior
- How metrics and algorithms can *create* human behavior (*reactivity*)
- The conflicts that can arise from the *differing worldviews* of data scientists and traditionalists (e.g., between epidemiologists and clinicians in health care, between data scientists and bench scientists in pharmaceutical research) and the potential loss of valuable organizational knowledge and skill that can result from decision automation
- How to redesign post-automation *jobs and control systems*

Environment

In addition, general managers need to know the broader environment of data analytics, a topic that is not generally well covered in Information Systems curricula now:

- Data ecosystems
  - Understand the complexity of the *data ecosystems* involved in, say, digital advertising, finance, or health care
  - Understand the business models of *platform operators* and the roles they play in the diffusion, governance, and dynamics of digital business
- Regulatory regime
  - Understand the broad outlines and complexity of the *international regulatory environment* as it relates to data and algorithms
  - Understand *compliance functions* in organizations, the roles of compliance specialists and *certification* programs, and the function of "*regtech*" (regulatory technologies)
  - Know that regulations *do not* cover all ethical responsibilities of organizations and their managers
  - Know that complying with regulations require *discretion and creative problem solving*

Of course, we could call this model TOPE, expanding the familiar Technology, Organization, and People acronym, but we prefer to call it ***POET*** instead, after the popular MBA website, Poets and Quants (https://poetsandquants.com/). We belief that POET offers a useful framework for developing curricula and educational materials in analytic ethics for general managers, and possibly for data scientists as well.

Conclusion

We have offered the POET framework as a starting point for debate about a proposed ethics education curriculum for general managers. Undoubtedly, it will undergo considerable revision over time. Regardless of its final shape, however, we still need to address how best to provide such education to emerging management leaders. We see at least three options: 1) via the basic data science course that most business and management programs have introduced (or will) for all their students, 2) through a general business ethics, business law, or corporate social responsibility course, or 3) some other option, yet to be devised.

The first two options each have drawbacks. It is rare enough for the topic of AI ethics to be adequately covered in entire data science *programs*, because instructors in those programs always claim that they cannot, as it is, cover all the essential *technical* material that students need. How likely is it, then, that something like our proposed curriculum would be given adequate treatment in a standalone course on business analytics for general managers? Similarly, the business ethics, business law and ethics, and corporate social responsibility courses we have reviewed generally—and appropriately—cover many issues unrelated to data and analytics. Furthermore, instructors of those courses typically have little background in the relevant Information Systems issues. So, finding good ways to address managers' needs for education in AI ethics is going to require lots of creative thinking—and the participation of Information Systems scholars and educations. We urge you to join us in this effort.

Acknowledgement

This research was supported in part by Bentley University's Data Innovation Network (BDIN), Heikki Topi, Principal Investigator.

References

Gentile, M. C. (2010). Giving voice to values: How to speak your mind when you know what? s Right, Yale University Press.

Jobin, A., et al. (2019). "The global landscape of AI ethics guidelines." Nature Machine Intelligence **1**(9): 389-399.

Markus, M. L. (2016). Obstacles on the Road to Corporate Data Responsibility. Big Data is Not a Monolith: Policies, Practices, and Problems. C. R. Sugimoto, H. R. Ekbia and M. Mattioli. Cambridge, MA, The MIT Press**:** 143-161.

Markus, M. L. (2017). "Datification, organizational strategy, and is research: what's the score?" The Journal of Strategic Information Systems **26**(3): 233-241.

Newell, S. and M. Marabelli (2015). "Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of 'datification'." The Journal of Strategic Information Systems **24**(1): 3-14.