
ETHICAL CODES AND LEARNING ANALYTICS

Stephen Downes, National Research Council Canada, Canada

Abstract

The growth and development of learning analytics has placed a range of new capacities into the hands of educational institutions. At the same time, this increased capacity has raised a range of ethical issues. A common approach to address these issues is to develop an ethical code of conduct for practitioners. Such codes of conduct are drawn from similar codes in other disciplines. Some authors assert that there are fundamental tenets common to all such codes. This paper consists of an analysis of ethical codes from other disciplines. It argues that while there is some overlap, there is no set of principles common to all disciplines. The ethics of learning analytics will therefore need to be developed on criteria specific to education. We conclude with some ideas about how this ethic will be determined and what it may look like.

Introduction

What distinguishes ethical codes from other forms of ethics generally is that while they may assign duties and responsibilities, these are assumed voluntarily by virtue of being a member of the profession. To become a nurse is, for example, to adopt as a personal code the ethical norms and values that define that particular profession.

The purpose of this chapter is to showcase the wide range of ethical codes that are employed in different professions, some of which are directly related to the use of analytics in that profession, and others which describe ethics in the profession generally. This diversity is not widely recognized; there is often a presumption, if not an explicit assertion, that the values in these ethical codes, and in ethics generally, are common, core, and universal.

This statement from Metcalf (2014) is typical: “There are several principles that can be found at the core of contemporary ethics codes across many domains:

- respect for persons (autonomy, privacy, informed consent),
- balancing of risk to individuals with benefit to society,
- careful selection of participants,

- independent review of research proposals,
- self-regulating communities of professionals,
- funding dependent on adherence to ethical standards.”

Whether or not one actually believes these principles are foundational, it remains a matter of empirical fact that they are not universal and not core. The same can be said for similar assertions of universality made elsewhere (for example: Pitofsky, (1998; p.7), Singer and Vinson (2002), CPA (2017)).

This chapter is a substantial survey of dozens of ethical codes. Though every attempt has been to keep this treatment brief, it is nonetheless not brief. By laying out the evidence I endeavour to show, rather than argue, that there is no common foundation to the ethical codes that govern different professions.

We'll begin with a quick overview of what we mean by ethical codes, discussing the purpose and operation of ethical codes, some of the components of ethical codes, and the ways in which these codes differ from each other. Then we'll take an extended look at the issues raised by the codes. First we look at what problems the codes are trying to solve, or in other words, what the purpose was for writing the codes. Then we look at a length list of values and priorities revealed in the codes. After this examination, we consider the question, to whom are the professionals described in the codes obligated? Finally, we ask what bases and foundations underlie the recommendations in the codes.

The full set of ethical codes is displayed, with readers invited to notice the ways in which they differ from each other, in Appendix 1: An Ethical Codes Reader, with references linking back to the full code in question, for further study as desired by the reader.

Standards of Conduct

Why Ethical Codes?

The need for professional ethics encompasses a number of factors. There is the need to be able to trust a person in a position of trust. There is the need to make good decisions and to do the right thing. And then there are various intangibles. The Project Management Institute (PMI, 2020) states, “Ethics is about making the best possible decisions concerning people, resources and the environment. Ethical choices diminish risk, advance positive results, increase trust, determine long term success and build reputations. Leadership is absolutely dependent on ethical choices.”

But these are not the only reasons advanced to justify professional ethics. There is the concern that without a statement of ethics, unethical conduct will abound. “The absence of a formal code could be seen almost as a guarantee that if such cases did exist they would

be swept under the carpet, left to others (probably the law) to sort out,” writes Sturges (2003).

Others are less concerned about good behaviour per se than they are about the bottom line. Alankar Karpe (2015), for example, writes in “Being Ethical is Profitable” that “Shortcuts and sleazy behaviour sometimes pay handsomely, but only for the short term. Organizations must remember that any benefits from lying, cheating, and stealing usually come at the expense of their reputation, brand image, and shareholders.” And, as he notes, “There is one and only one social responsibility of business – to use it[s]resources and engage in activities designed to increase its profits so long as it stays within the rules of the game, which is to say, engages in open and free competition.”

Additionally, there are services and institutions that require professional ethics in order to function. For example, the CFA Institute (2017) states, “ethical conduct is vital to the ongoing viability of the capital markets.” It notes, “compliance with regulation alone is insufficient to fully earn investor trust. Individuals and firms must develop a ‘culture of integrity’ that permeates all levels of operations.” Indeed, it is arguable that society as a whole could not function without professional ethics. Thus, the “CFA Institute recently added the concept ‘for the ultimate benefit of society’ to its mission.”

Certain disciplines see ethical codes as essential to being recognized as a profession. Hence, for example, for librarians, “Keith Lawry set the idea of a code in a particularly positive view of the professionalization process in British librarianship. He linked the Library Association’s possession of a code of professional conduct with the potential for statutory recognition of the association’s control of who might and who might not practise librarianship” (Sturges, 2003)

Finally, practitioners need them. As Rumman Chowdhury, Accenture’s Responsible AI Lead, said, “I’ve seen many ‘ethics codes’ focused on AI, and while many of them are very good they’re more directional than prescriptive – more in the spirit of the Hippocratic Oath that doctors are expected to live by. Meanwhile, many data scientists are hungry for something more specific and technical. That’s what we need to be moving toward” (De Bruijn, et al., 2019)

Ethical Codes as Standards of Conduct

While ethics commonly applies to people in general, there is a specific class of ethics that applies to people by virtue of their membership in a professional group. There are different approaches, but in general, “professional ethics are principles that govern the behaviour of a person or group in a business environment. Like values, professional ethics provide rules

on how a person should act towards other people and institutions in such an environment” (Government of New Zealand, 2018).

Professional ethics can be characterized as imposing a higher standard of conduct. The reasons for this vary, but (as we discuss below) a higher standard is demanded because professionals are in positions of power, they have people in their care, and they are expected to have special competencies and responsibilities. Additionally, professional ethics may require that practitioners put the interests of others ahead of their own. This may include duties not only to those in one’s care, but also to clients, organizations, or even intangibles like “the Constitution” or “the public good”.

As such, professional ethics are often expressed in terms of codes of conduct (indeed, it is hard to find a sense of professional ethics where such a code is not employed). Though the code is normative (“breaches of a code of conduct usually do carry a professional disciplinary consequence” (Ibid.)) usually the intent of the code is to remind professionals of their duty and prompt them regarding specific obligations.

Ethical Codes as Requirements

In the world of software engineering, in addition to ethical standards as codes of conduct, ethical codes can be seen as defining requirements. This is proposed, for example, by Guizzardi et al. (2020), they write, “Ethical requirements are requirements for AI systems derived from ethical principles or ethical codes (norms). They are akin to Legal Requirements, i.e., requirements derived from laws and regulations.” Ethical requirements are drawn from stakeholders in the form of principles and codes. From these, specific requirement statements are derived. “For example, from the Principle of Autonomy one may derive “Respect for a person’s privacy”, and from that an ethical requirement “Take a photo of someone only after her consent” (Ibid; p.252).

An important distinction between the idea of ethical codes as standards of conduct and ethical codes as requirements is that in the former case, the AI is treated as an ethical agent can reason and act on the basis of ethical principle, while in the latter case, the AI is not. “Rather, they are software systems that have the functionality and qualities to meet ethical requirements, in addition to other requirements they are meant to fulfill” (Ibid; p.252).

As Opposed to Legal Requirements

We stated above that “ethics is not the same as the law”. This is a case where that principle applies. What we are interested in here is the sense of an ethical code as a principle of ethics, not as a legal document. It reflects the fact that a person chooses a profession for themselves, and thereby voluntarily enters into a set of obligations characterized by that

profession. “Professions must be ‘professed’ (that is, declared or claimed)” (Davis, 2010; p.232).

Thus we may say that ethics may be influenced by, but are distinct from, the following (all from Government of New Zealand, 2018):

- Fiduciary duties – fiduciary duties are “special obligations between one party, often with power or the ability to exercise discretion that impacts on the other party, who may be vulnerable” (Wagner Sidlofsky, 2020). Examples of fiduciary relations include those between lawyer and client, trustee and beneficiary, director and company, power of attorney and beneficiary and accountant and client.
- Contractual obligations – these require the professional to perform the terms of the contract, and “includes a duty to act with diligence, due care and skill, and also implies obligations such as confidentiality and honesty” (New Zealand, 2018).
- Other laws – for example, In New Zealand this could include the Consumer Guarantees Act 1993.

What distinguishes legal requirements, arguably, from ethical principles is the element of choice. In the case of legal requirements, the law compels you to behave in a certain way, with increasing penalties for non-compliance. In an important sense, it doesn’t matter whether the law or the principle in question is ethical or not. You are penalized if you do not comply.

It may be argued that the relation between ethics and law is such that in a treatment of the ethics of learning analytics we ought also to be concerned with the law in relation to learning analytics. We will see this come up in two ways: first, in the argument that “obeying the law” is part of the ethical responsibility of a practitioner, and second, in the argument that the law regarding learning analytics is or ought to be informed by ethical principles.

Principles and Values

“Values are general moral obligations while principles are the ethical conditions or behaviors we expect” (Gilman, 2005; p.10). Values and principles are connected. As Terry Cooper (1998:12) explains, “An ethical principle is a statement concerning the conduct or state of being that is required for the fulfillment of a value; it explicitly links a value with a general mode of action.” For example, we may state that we value “justice”, but we would need a principle like “treat equals equally and unequals unequally” to explain what we mean by “justice”.

All ethics codes encompass both principles and values, though (as we shall see below) usually more implicitly than explicitly. Values (such as honesty and trustworthiness) are

often assumed tacitly, as not needing to be stated. Sometimes they are expressed in a preamble to the code, not as an explicit list, but rather in the sense of establishing a context. For example, the Canadian Code of Public Service ethical code has a preamble describing the role of the public service, as well as a listing of the fundamental values (TBS, 2011).

The Value of Professional Codes

Codes of professional ethics or conduct are widely used. They bring a utilitarian value to the conversation. They provide a framework for professionals carrying out their responsibilities. They clearly articulate unacceptable conduct. And they provide a vision toward which a professional may be striving (Gilman, 2005; p.5). Having a code, it is argued, is key to the prevention of unacceptable conduct. That's why, for example, the United Nations Convention Against Corruption included a public service code of conduct as an essential element in corruption prevention, says Gilman (Ibid). Yet the convention is an interesting example: there is no code of conduct for the private sector. Why?

At the same time, it is argued that "Codes are not designed for 'bad' people, but for the persons who want to act ethically" (Ibid; p.7). That is, they provide guidance for a person who wants to act ethically, but who may not know what is right. Therefore, codes are preventative only in the sense that they prevent conduct that is accidentally unacceptable. They may seem to be unnecessary in the case of a well-developed profession and body of professionals, but in a new environment, such as data analytics in education, there is much that is not yet clearly and widely understood.

Moreover, argues Gilman, a code of ethics will change the behaviour of bad actors, even if it does not incline them toward good. "When everyone clearly knows the ethical standards of an organization they are more likely to recognize wrongdoing; and do something about it. Second, miscreants are often hesitant to commit an unethical act if they believe that everyone else around them knows it is wrong. And, finally corrupt individuals believe that they are more likely to get caught in environments that emphasize ethical behavior." (Ibid; p.8)

Study of Ethical Codes

More than 70 ethical codes were studied as a part of this review. The selection methodology undertaken was designed to encourage as wide a range of ethical codes as possible. To begin, ethical codes referenced in relevant metastudies (such as) were evaluated. Codes referenced by these ethical codes were studied, to establish a history of code development within a discipline. Documents from relevant disciplinary associations were studied, to find more ethical codes. The selection of ethical codes includes the following major disciplinary groups (and the number of individual codes studied).

Ethical Codes and Learning Analytics

- Professional ethics – broad-based ethical codes (4);
- Academic ethics – codes of conduct for professors and staff in traditional academic institutions (3);
- Teacher ethics – codes governing teachers and the teaching profession (7);
- Ethics for librarians and information workers – ethics of information management (2);
- Public service ethics – codes of conduct for government employees (2);
- Research ethics – includes international declarations and government policy (6);
- Health care ethics – including codes for doctors and nurses (6);
- Ethics in social science research – research ethics (1);
- Data ethics – government and industry declarations on the use of study and survey data (7);
- Market research ethics – codes describing the ethical use of data in advertising and market studies (2);
- Journalism ethics – codes of conduct governing the use of public information by journalists (3);
- Ethics for IT professionals – system administration and software development ethics (3);
- Data research ethics – related specifically to the use of data in research (1);
- Ethics for artificial intelligence – government, industry and academic codes (15);
- Information and privacy – principles specifically addressing individual rights (1);
- Ethics in educational research – policies governing educational researchers specifically (3);
- Ethics in learning analytics – government, academic and industry guidelines and codes (7).

How the Codes Differ

Metcalf (2014) identifies a number of the reasons ethical codes vary across professions, and even within professions (quotes in the list below are all from Metcalf):

- Motivation: The events that prompt the development of ethical codes; for example, “in biomedicine, ethics codes and policies have tended to follow scandals” while by contrast “major policies in computing ethics have presaged many of the issues that are now experienced as more urgent in the context of big data.”
- Purpose: “Analyses of ethics codes note a wide range of purposes for ethics codes (Frankel, 1998; Gaumintz and Lere, 2002; Kaptein and Wempe, 1998).”
- Interests: “Frankel (1989) notes that all ethics codes serve multiple interests and therefore have multiple, sometimes conflicting, dimensions. He offers a taxonomy of aspirational, educational, and regulatory codes.”

Ethical Codes and Learning Analytics

- Burden: who does the ethical code apply to? Metcalf notes that “greater burdens are placed on individual members to carry out the profession’s ethical agenda,” but different burdens may fall on different groups of people.
- Enforcement: “Organizations, institutions and communities tend to develop methods of enforcement that reflect their mission.”

Each code of ethics was subjected to an analysis that includes the following criteria:

- What ethical issues is it attempting to address (for example, is focused on malpractice, on conflict of interest, on violation of individual rights, etc.)?
- What are its core values or highest priorities (as opposed to the detailed specification of ethical principles described, as defined by Cooper (1998; p.12), Gilman (2005; p.10))?
- Which ethical issues from the literature of learning analytics issues do they address?
- Who is governed, and to whom are they obligated? (e.g., AITP (2017) list six separate groups to which information professionals have obligations).
- What is the basis (if any) for the statement of ethical values and principles (for example, the Royal Society’s recommendations are based in a “public consultation” (Drew, 2018)), while numerous other statements are based in principles such as “fairness” and “do no harm”.

Applications of Learning Analytics

Analytics is thought of generally as “the science of examining data to draw conclusions and, when used in decision making, to present paths or courses of action.” (Picciano, 2012). This includes not only the collection of the data but also the methods of preparation and examination employed, and the application of the data in decision-making. Thus the term *analytics* can be thought of as the overall process of “developing actionable insights through problem definition and the application of statistical models and analysis against existing and/or simulated future data” (Cooper, 2012).

The focus of this paper is the use of analytics as applied to learning and education (typically called *learning analytics*). Learning analytics is typically defined in terms of its objective, which is to improve the chance of student success (Gasevic, Dawson, & Siemens, 2015). Accordingly, when founding the Society for Learning Analytics (SoLAR) George Siemens defined learning analytics as “the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimizing learning and the environments in which it occurs” (Siemens, 2012).

We apply a broad definition of learning analytics. A wider definition not only avoids the difficulties of establishing a more narrow definition, but also ensures we do not disregard potential ethical implications simply because the practice is “outside the scope of learning

analytics”. Arguing for a broader definition of analytics necessarily leads us to consider including artificial intelligence (AI) in the conversation. However you define the terms, artificial intelligence plays a significant role in analytics, and vice versa, so we will treat them together as one thing (Adobe Experience Cloud Team, 2018). If a distinction is necessary during the course of the discussion, we will apply it.

Potential applications of learning analytics are based on what analytics can do and how they work. Modern analytics is based mostly in machine learning and neural networks, and these in turn provide algorithms for pattern recognition, regression, and clustering. Built on these basic capabilities are four widely-used categories (Brodsky et al., 2015; Boyer & Bonnin, 2017) to which we add additional fifth and sixth categories, generative analytics and deontic analytics:

- descriptive analytics, answering the question “what happened?”;
- diagnostic analytics, answering the question “why did it happen?”;
- predictive analytics, answering the question “what will happen?”;
- prescriptive analytics, answering the question “how can we make it happen?”; and
- generative analytics, which use data to create new things, and
- deontic analytics, answering the question “what should happen?”.

Within each of these categories we can locate the various applications that fall under the heading “learning analytics”.

Descriptive Analytics

Descriptive analytics include analytics focused on description, detection and reporting, including mechanisms to pull data from multiple sources, filter it, and combine it. The output of descriptive analytics includes visualizations such as pie charts, tables, bar charts or line graphs. Descriptive analytics can be used to define key metrics, identify data needs, define data management practices, prepare data for analysis, and present data to a viewer. (Vesset, 2018). Tracking is an important part of descriptive analytics. The purpose of tracking is to measure systems performance and institutional compliance. Relative costs and benefits are compared to find the most cost-effective solution (Ware et al., 1973; p.9).

Higher education institutions also use descriptive analytics to construct student profiles. A person’s learning activities, for example, can be graphed and displayed in comparison with other learners. This analysis can contain fine-grained detail, for example, attention metadata (Duval, 2011). Today, a standardized format, the Experience API, is used to collect and store activity data in a Learning Record Store (LRS) (Corbí & Solans, 2014; Kevan & Ryan, 2016). These support dashboards such as LAViEW (Learning Analytics Visualizations & Evidence Widgets) that helps learners analyse learning logs and provide evidence of learning. (Majumdar et al., 2019)

Similar functionality is also provided by IMS Global's Caliper learning analytics (Oakleaf et al., 2017)

Diagnostic Analytics

Diagnostic analytics look more deeply into data in order to detect patterns and trends. Such a system could be thought of as being used to draw an inference about a piece of data based on the patterns detected in sample or training data, for example, to perform recognition, classification or categorization tasks. Diagnostic analytics are applied in a wide range of applications.

Security applications are common. To support physical security, facial and object recognition technology is being used in schools and institutions. For example, a New York school district is using an application called AEGIS to identify potential threats (Klein, 2020). For digital security, analytics applications that help filter unwanted messages (whether sent by humans or bots) are generally available and widely used. Users can learn to train their own machine learning to filter spam (Gan, 2018) or use commercial systems such as Akismet (Barron, 2018). Automated fakes detection systems are becoming more widely used (Li & Lyu, 2019).

Diagnostic analytics is also employed to ensure academic discipline. Pattern recognition, for example, is used for plagiarism detection Amigud et al. (2017). Video recognition and biometrics are also used for security purposes and exam proctoring (Rodchua, 2017). "For instance, Examity also uses AI to verify students' identities, analyze their keystrokes, and, of course, ensure they're not cheating. Proctorio uses artificial intelligence to conduct gaze detection, which tracks whether a student is looking away from their screens" (Heilweil, 2020).

There is a large literature devoted to automated grading, beginning with Page (1966), continuing through the Hewlett competition (Kaggle, 2012), and today the technology has at least "developed to the point where the systems provide meaningful feedback on students' writing and represent a useful complement (not replacement) to human scoring" (Kaja & Bosnic, 2015). Ultimately, AI could replace grading altogether. Rose Luckin argues, "logging every keystroke, knowledge point and facial twitch, then the perfect record of their abilities on file could make such testing obsolete" (Beard, 2020). This creates the possibility of assessing competencies from actual performance data outside educational environments, for example, using technologies like analytics-based assessment of personal portfolios (van der Schaaf et al., 2017) or using data-driven skills assessment in the workplace (Lin et al., 2018).

Predictive Analytics

Numerous products and studies are based on the idea that “analytics tools can identify factors statistically correlated with students at risk of failing or dropping out.” (Scholes, 2016; Gasevic, Dawson, & Siemens, 2015). For example, a Jisc report describes several such projects, including one at New York Institute of Technology (NYIT) that used four data sources: “admission application data, registration / placement test data, a survey completed by all students, and financial data” (Sclater, Peasgood, and Mullan, 2016). Student retention is also supported by predictive analytics. Predictive analytics is also used to assist in learning design, including adaptive learning design. “Findings indicated that the primary predictor of academic retention was how teachers designed their modules, in particular the relative amount of so-called ‘communication activities’.” (Rientes & Jones, 2019; p.116)

Analytics can also draw from campus information sources to support student advising. For example, the Berkeley Online Advising (BOA, 2020) project at the University of California at Berkeley “integrates analytical insights with relationship and planning tools for advisors of large cohorts and the students they support” (Heyer & Kaskiris, 2020). Additionally, the Comprehensive Analytics for Student Success (COMPASS) project at the University of California, Irvine, “focuses on bringing relevant student data to campus advisors, faculty, and administrators... to improve undergraduate student outcomes” (UCI Compass, 2020). As O’Brien (2020) writes, “These tools provide advisors with information that allows for proactive outreach and intervention when critical student outcomes are not met.” Combining these approaches is an initiative called “precision education”. Yang and Ogata (2020) suggest that analogous to precision medicine, precision education systems consider a wider array of variables than learning analytics, “students’ IQ, learning styles, learning environments, and learning strategies.”

Prescriptive Analytics

An oft-cited application is the potential of learning analytics to make content recommendations, either as a starting point, or as part of a wider learning analytics-supported learning path. For example, the Personalised Adaptive Study Success (PASS) system supports personalisation for students at Open Universities Australia (OUA) (Sclater, Peasgood, & Mullan, 2016). Students report desiring recommendations regarding potential learning activities, and suggestions for potential learning partners. (Schumacher, 2018) Content and learning path recommendations are based not only on the discipline being studied but also on the individual learning profile, academic history, and a variety of contextual factors. (Ifenthaler & Widanapathirana, 2014)

Adaptive learning is a step beyond learning recommendations in the sense that the learning environment itself changes (or “adapts”) to events in the learning experience (Sonwalkar, 2007). For example, “Adaptive learning systems – like IBM Watson and Microsoft Power BI – have the advantage of continually assessing college students’ skill and confidence levels.” (Neelakantan, 2019). Early adaptive learning applications were expert systems based on explicit knowledge representations and user models, that is, they were based on statements and rules (Garrett & Roberts, 2004). More recently, the “black box” methods characteristic of contemporary analytics, such as neural networks, have been employed (Almohammadi et al., 2017).

Generative Analytics

Generative analytics is different from the previous four categories in the sense that it is not limited to answering questions like “what happened” or “how can we make it happen”, but instead uses the data to create something that is genuinely new. In a sense, it is like predictive and prescriptive analytics in that it extrapolates beyond the data provided, but while in the former two we rely on human agency to act on the analytics, in the case of generative analytics the analytics engine takes this action on its own.

In addition to emulating human conversation, chatbots will generate additional human responses, such as gestures and emotions. For example, there’s Magic Leap’s Mica, an AI-driven being “that comes across as very human” (Craig, 2018). “What is remarkable about Mica is not the AI, but the human gestures and reactions (even if they are driven by AI).” Meanwhile, though “fictionalized and simulated for illustrative purposes only”, products like Samsung’s Neon are being called “artificial humans”, “a computationally created virtual being that looks and behaves like a real human, with the ability to show emotions and intelligence.” (Craig, 2020)

Analytics engines, provided with data, can generate content. The Washington Post uses an AI called Heliograf to write news and sports articles; in its first year it wrote around 850 items. “That included 500 articles around the election that generated more than 500,000 clicks.” (Moses, 2017) Analytics and AI have self-generated computer science papers (Stribling et al., 2005), music (Galeon, 2016), art (Shepherd, 2016), books (Springer Nature, 2019) and inventions (Fleming, 2018). There are now commercial AI-based applications that generate educational resources, including articles (e.g., AiWriter), textbooks, test questions (e.g. WeBuildLearning), and more.

Such technology can make educational content more interesting and engaging. For example, In 2015, an algorithm called DeepStereo developed for Google Maps was able to generate a video from a series of still photographs (Flynn et al., 2015). Also, “With deep fakes, it will be possible to manufacture videos of historical figures speaking directly to

students, giving an otherwise unappealing lecture a new lease on life” (Chesney & Citron, 2018; p.1769). Chesney and Citron write, “The educational value of deep fakes will extend beyond the classroom. In the spring of 2018, BuzzFeed provided an apt example when it circulated a video that appeared to feature Barack Obama warning of the dangers of deep-fake technology itself. One can imagine deep fakes deployed to support educational campaigns by public-interest organizations such as Mothers Against Drunk Driving (Chesney & Citron, 2018; p.1770).

It may seem far-fetched, but some pundits are already predicting the development of artificial intelligences and robots teaching in the classroom. In a recent celebrated case, a professor fooled his students with Jill Watson, an artificial tutor (Miller, 2016). “Yuki’, the first robot lecturer, was introduced in Germany in 2019 and has already started delivering lectures to university students at The Philipps University of Marburg.” (Ameen, 2019). While most observers still expect AI and analytics to be limited to a support role, these examples suggest that the role of artificial teachers might be wider than expected.

Deontic Analytics

There is an additional question that needs to be answered, and has been increasingly entrusted to analytics: “what ought to happen?” Recently the question has been asked with respect to self-driving vehicles in the context of Philippa Foot’s “trolley problem”. (Foot, 1967). In a nutshell, this problem forces the reader to decide whether to take an action to save six and kill one, or to desist from action to save one, allowing (by inaction) six to be killed. It is argued that automated vehicles will face similar problems.

It may be argued that these outcomes are defined ahead of time by human programmers. But automated systems have an impact on what content is acceptable (and what is not) in a society. We see this in effect on online video services. “On both YouTube and YouTube Kids, machine learning algorithms are used to both recommend and mediate the appropriateness of content” (UC Berkeley Human Rights Center Research, 2019). Though such algorithms are influenced by input parameters, their decisions are always more nuanced than designed, leading people to adapt to the algorithm, thereby redefining what is acceptable.

What counts as “appropriate” behaviour may be shaped by analytics and AI. These and additional implications are being investigated by HUMAINT, “an interdisciplinary JRC project aiming to understand the impact of machine intelligence on human behaviour, with a focus on cognitive and socio-emotional capabilities and decision making” (Tuomi, 2018; HUMAINT, 2020). An AI can select between what might be called “good” content and “bad” content, displaying a preference for the former. For example, in response to violence in conflict zones, researchers “argue the importance of automatic identification

of user-generated web content that can diffuse hostility and address this prediction task, dubbed ‘hope-speech detection’” (Palakodety et al., 2019).

There is another line of research that proposes that AI can define what’s fair. An early example of this is software designed to optimize the design of congressional voting districts in such a way that minimizes gerrymandering (Cohen-Addad, Klein, & Young, 2018). In another study, research suggested that “an AI can simulate an economy millions of times to create fairer tax policy” (Heaven, 2020). A tool developed by researchers at Salesforce “uses reinforcement learning to identify optimal tax policies for a simulated economy.” The idea in this case was to find tax policy that maximized productivity and income equality in a model economy.

It should be noted that discussion of generative and deontic analytics lie outside most traditional accounts of analytics and ethics. And it is precisely in these wider accounts of analytics that our relatively narrow statements of ethical principles are lacking. It is possible to apply analytics correctly and yet still reach a conclusion that would violate our moral sense. And it is possible to use analytics correctly and still do social and cultural harm. An understanding of ethics and analytics may begin with ethical principles, but it is far from ended there.

Ethical Issues in Learning Analytics

We will follow Narayan (2019), who classifies the ethical issues in learning analytics under three headings: issues that arise when analytics works, issues that arise because analytics are not yet reliable, and issues that arise in cases where the use of analytics seems fundamentally wrong. To these three sets of issues we will add a fourth describing wider social and cultural issues that arise with the use of analytics and AI, and a set of issues related specifically to bad actors.

Many of these issues will be familiar to readers, for example, the potential misuse of facial recognition, surveillance and tracking, AI-based assessment, misrepresentation and prejudice, explainability, filter bubbles and feedback effects. Others are less frequently discussed but raised equally serious ethical issues, for example, the mechanisms for appealing AI-based evaluations, systems consistency reliability, stalking, alienation, network effects (i.e., winner takes all), and environmental impact.

When Analytics Works

Modern AI and analytics work. As Mark Liberman (2019) observes, “Modern AI (almost) works because of machine learning techniques that find patterns in training data, rather than relying on human programming of explicit rules.” This is in sharp contrast to earlier rule-based approaches that “generally never even got off the ground at all.”

Analytics and AI require data above all, and so in order to support this need institutions and industries often depend on surveillance. However, “when in wrong hands, these systems can violate civil liberties.” (UC Berkeley, 2019) Once surveillance becomes normal, its use expands (Marx, 2020). Private actors, as well, employ surveillance for their own purposes. For example, Amazon-owned Whole Foods is tracking its employees with a heat map tool that ranks stores most at risk of unionizing (Peterson, 2020). Analytics makes tracking accessible to everyone. “Miniature surveillance drones, unseen digital recognition systems, and surreptitious geolocational monitoring are readily available, making long-term surveillance relatively easy and cheap” (Cavoukian, 2013; p.23).

Analytics also erodes our ability to be anonymous. This is partially because of spying and tracking, and partially because data about individuals can be cross-referenced. “When Facebook acts as a third-party tracker, they can know your identity as long as you’ve created a Facebook account and are logged in – and perhaps even if you aren’t logged in” (Princiya, 2018). And analytics arguably creates a social need to eliminate anonymity. As Bodle argues, “A consensus is growing among governments and entertainment companies about the mutual benefits of tracking people online.” Hence, provisions against anonymity, he argues, are being built into things like trade agreements and contracts.

Recent debate has focused on the use of facial recognition technologies, with IBM, Microsoft and Amazon all announcing they will cease efforts. A startup called Clearview AI makes the risk clear. “What if a stranger could snap your picture on the sidewalk then use an app to quickly discover your name, address and other details? has made that possible” (Moyer, 2020). Mark Andrejevic and Neil Selwyn (2019) outline a number of additional ethical concerns involving facial recognition technology in schools: the dehumanising nature of facially focused schooling, the foregrounding of students’ gender and race, the increased authoritarian nature of schooling, and more.

The previous sections each raise their own issues, but all touch on the issue of privacy generally. While it may be argued that privacy protects the powerful, at the expense of the weaker (Shelton (2017), “Personal privacy is about more than secrecy and confidentiality. Privacy is about being left alone by the state and not being liable to be called to account for anything and everything one does, says or thinks” (Cavoukian, 2013; p.18). We might say people should be able to live their lives in “quiet enjoyment” of their possessions, property and relationships (Andresi, 2019).

In education learning analytics used for assessment can score student work with accuracy and precision. Students recognize this. But students have mixed feelings about such systems, preferring “comments from teachers or peers rather than computers.” (Roscoe et al., 2017) It is arguable that students may prefer human assessment because they may

feel more likely to be seen as an individual with individual flair, rather than erroneously deviating from the expectations of the analytics engine. As one college official says, “Everyone makes snap judgments on students, on applicants, when first meeting them. But what worries me about AI is AI can’t tell the heart of a person and the drive a person has.”

Humans often use discretion when applying the rules. “Organizational actors establish and re-negotiate trust under messy and uncertain analytic conditions” (Passi & Jackson, 2018) In the case of learning analytics, Zeide (2019) writes that a human instructor might overlook a student’s error “if she notices, for example, that the student clearly has a bad cold.” By contrast, “Tools that collect information, particularly based on online interactions, don’t always grasp the nuances.” The impact of a lack of discretion is magnified by uncertainties in the data that might be recognized by a human but overlooked by the machine (Passi & Jackson, 2018; Malouff & Thorsteinsson, 2016). There is a need for a principle of “remedy for automated decision” that is “fundamentally a recognition that as AI technology is deployed in increasingly critical contexts, its decisions will have real consequences, and that remedies should be available just as they are for the consequences of human actions” (Fjeld et al., 2020; p.33).

Analytics can also be used to create misleading images and videos (Chesney & Citron, 2018; p.1760) write “To take a prominent example, researchers at the University of Washington have created a neural network tool that alters videos so speakers say something different from what they originally said.” There are numerous unethical uses of content manipulation, including exploitation, sabotage, harm to society, distortion of discourse, manipulation of elections, erosion of trust, exacerbation of divisions, undermining of public safety, and undermining journalism (Ibid; pp.1772-1786).

A number of recent high-profile cases have raised the possibility of analytics being used to (illegitimately?) manipulate the thoughts, feelings and emotions of users. For example, one study experimented on Facebook users (without their knowledge or consent) to show that “emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness” (Kramer, Guillory, & Hancock, 2014). An article from RAND suggests, “Whoever is first to develop and employ such systems could easily prey on wide swaths of the public for years to come” (Paul & Posard, 2020).

Manipulation of the user can be used for beneficial purposes, as described above. However it becomes ethically problematic when the institution, rather than the user, benefits. As Kleber (2018) writes, “Casual applications like Microsoft’s XiaoIce, Google Assistant, or Amazon’s Alexa use social and emotional cues for a less altruistic purpose – their aim is to

secure users' loyalty by acting like new AI BFFs. Futurist Richard van Hooijdonk quips: "If a marketer can get you to cry, he can get you to buy." Moreover, Kleber continues, "The discussion around addictive technology is starting to examine the intentions behind voice assistants. What does it mean for users if personal assistants are hooked up to advertisers? In a leaked Facebook memo, for example, the social media company boasted to advertisers that it could detect, and subsequently target, teens' feelings of 'worthlessness' and 'insecurity', among other emotions (Levin, 2017)".

Schneier (2020) writes, "The point is that it doesn't matter which technology is used to identify people... The whole purpose of this process is for companies – and governments – to treat individuals differently." In many cases, differential treatment is acceptable. However, in many other cases, it becomes subject to ethical concerns. The accuracy of analytics creates an advantage for companies in a way that is arguably unfair to consumers. For example, the use of analytics data to adjust health insurance rates (Davenport & Harris, 2007) works in favour of insurance companies, and thereby, arguably, to the disadvantage of their customers. Analytics are used similarly in academics, sometimes before the fact, and sometimes after. For example, in a case where failure was determined by predicted learning events, the "Mount St. Mary's University... president used a survey tool to predict which freshman wouldn't be successful in college and kicked them out to improve retention rates" (Foresman, 2020).

When it doesn't

Artificial Intelligence and analytics often work and as we've seen above can produce significant benefits. On the other hand, as Liberman comments (2019), AI is brittle. When the data are limited or unrepresentative, it can fail to respond to contextual factors our outlier events. It can contain and replicate errors, be unreliable, be misrepresented, or even defrauded. In the case of learning analytics, the results can range from poor performance, bad pedagogy, untrustworthy recommendations, or (perhaps worst of all) nothing at all.

Analytics can fail because of error, and this raises ethical concerns. "Analytics results are always based on the data available and the outputs and predictions obtained may be imperfect or incorrect. Questions arise about who is responsible for the consequences of an error, which may include ineffective or misdirected educational interventions" (Griffiths et al., 2016; p.4).

Analytics requires reliable data, "as distinguished from suspicion, rumor, gossip, or other unreliable evidence" (Emory University Libraries, 2019). Meanwhile, a "reliable" system of analytics is one without error and which can be predicted to perform consistently, or in other words, "an AI experiment ought to 'exhibit the same behavior when repeated under the same conditions' and provide sufficient detail about its operations that it may be

validated (Fjeld et al., 2020; p.29; Slade & Tait, 2019). Both amount to a requirement of “verifiability and replicability” of both data and process.

Additionally, the reliability of models and algorithms used in analytics “concerns the capacity of the models to avoid failures or malfunction, either because of edge cases or because of malicious intentions. The main vulnerabilities of AI models have to be identified, and technical solutions have to be implemented to make sure that autonomous systems will not fail or be manipulated by an adversary” (Hamon, Junklewitz, & Sanchez, 2020; p.2). But it is not yet clear that learning analytics are reliable (Contact North, 2018). For example, inconsistency can magnify ethical issues, especially in real-time analytics. “‘When the facts change, I change my mind’ can be a reasonable defence: but in order to avoid less defensible forms of inconsistency, changing your mind about one thing may require changing it about others also” (Boyd, 2019).

Additionally, there are widespread concerns about bias in analytics. In one sense, it is merely a specific way analytics can be in error or unreliable. But more broadly, the problem of bias pervades analytics: it may be in the data, in the collection of the data, in the management of the data, in the analysis, and in the application of the analysis. The outcome of bias is reflected in misrepresentation and prejudice. For example, “the AI system was more likely to associate European American names with pleasant words such as ‘gift’ or ‘happy’, while African American names were more commonly associated with unpleasant words.” (Devlin, 2017) “The tales of bias are legion: online ads that show men higher-paying jobs; delivery services that skip poor neighborhoods; facial recognition systems that fail people of color; recruitment tools that invisibly filter out women” (Powles & Nissenbaum, 2018).

Another source of error is misinterpretation. Because analytical engines don’t actually know what they are watching, they may see one thing and interpret it as something else. For example, looking someone in the eyes is taken as a sign that they are paying attention. And so that’s how an AI interprets someone looking straight at it. But it might just be the result of a student fooling the system. For example, students being interviewed by AI are told to “raise their laptop to be eye level with the camera so it appears they’re maintaining eye contact, even though there isn’t a human on the other side of the lens” (Metz, 2020). The result is that the AI misinterprets laptop placement as “paying attention”.

There is a risk, writes Ilkka Tuomi (2018), “that AI might be used to scale up bad pedagogical practices. If AI is the new electricity, it will have a broad impact in society, economy, and education, but it needs to be treated with care.” For example, badly constructed analytics may lead to evaluation errors. “Evaluation can be ineffective and even harmful if naively done ‘by rule’ rather than ‘by thought’” (Dringus, 2012). Even more

concerning is how poorly designed analytics could result in poorly defined pedagogy. Citing Bowker and Star (1999), Buckingham Shum and Deakin Crick (2012) argue that “a marker of the health of the learning analytics field will be the quality of debate around what the technology renders visible and leaves invisible, and the pedagogical implications of design decisions.”

Social and Cultural Issues

This is a class of issues that addresses the social and cultural infrastructure that builds up around analytics. These are not issues with analytics itself, but with the way analytics changes our society, our culture, and the way we learn.

Analytics is ethically problematic in society when it is not transparent. When a decision-making system is opaque, it is not possible to evaluate whether it is making the right decision. You might not even know the decision was made by a machine. Analytics requires a “principle of notification” (Fjeld et al., 2020; p.45). Additionally, transparency applies to the model or algorithm applied in analytics. “Transparency of models: it relates to the documentation of the AI processing chain, including the technical principles of the model, and the description of the data used for the conception of the model. This also encompasses elements that provide a good understanding of the model, and related to the interpretability and explainability of models” (Hamon, Junklewitz, & Sanchez, 2020; p.2).

Explainability is closely related to transparency. In the case of analytics, explainability seems to be inherently difficult. We’re not sure whether we’ll be able to provide explanations. Zeide (2019) writes, “Unpacking what is occurring within AI systems is very difficult because they are dealing with so many variables at such a complex level. The whole point is to have computers do things that are not possible for human cognition.” As Eckersley et al. (2017) say, “Providing good explanations of what machine learning systems are doing is an open research question; in cases where those systems are complex neural networks, we don’t yet know what the trade-offs between accurate prediction and accurate explanation of predictions will look like.”

Numerous agencies have announced efforts to ensure that automated decisions are “accountable” (Rieke, Bogen, & Robinson, 2018). But the nature of AI might make accountability impossible. “Suppose every single mortgage applicant of a given race is denied their loan, but the Machine Learning engine driving that decision is structured in such a way that the relevant engineers know exactly which features are driving such classifications. Further suppose that none of these are race-related. What is the company to do at this point?” (Danzig, 2020).

What we don't know might hurt us. The UK House of Lords Select Committee notes that "The use of sophisticated data analytics for increasingly targeted political campaigns has attracted considerable attention in recent years, and a number of our witnesses were particularly concerned about the possible use of AI for turbo-charging this approach" (Clement-Jones et al, 2018; para260). One example is the use of bot Twitter accounts to sow division during the Covid-19 pandemic. "More than 100 types of inaccurate COVID-19 stories have been identified, such as those about potential cures. But bots are also dominating conversations about ending stay-at-home orders and 'reopening America,' according to a report from Carnegie Mellon (Young, 2020)".

An ethical issue here arises because "information is filtered before reaching the user, and this occurs silently. The criteria on which filtering occurs are unknown; the personalization algorithms are not transparent" (Bozdag & Timmermans, 2011). Additionally, "We have different identities, depending on the context, which is ignored by the current personalization algorithms" (Ibid). Moreover, algorithms that drive filter bubbles may be influenced by ideological or commercial considerations (Introna & Nissenbaum, 2000; p.177). The eventual consequence may be disengagement and alienation. "Will Hayter, Project Director of the Competition and Markets Authority, agreed: '... the pessimistic scenario is that the technology makes things difficult to navigate and makes the market more opaque, and perhaps consumers lose trust and disengage from markets'" (Clement-Jones et al, 2018; para52).

Artificial intelligence and analytics impose themselves as a barrier between one person and another, or between one person and necessary access to jobs, services, and other social, economic and cultural needs. Consider the case of a person applying for work where analytics-enabled job applicant screening is being used. However, "La difficulté, pour les candidats pris dans les rets de ces systèmes de tris automatisés, est d'en sortir, c'est-à-dire se battre contre les bots, ces gardiens algorithmiques, pour atteindre une personne réelle capable de décider de son sort (The difficulty for candidates caught in the nets of these automated sorting systems is to get out of them, that is, to fight against bots, those algorithmic guardians, to reach a real person capable of deciding on their exit)" (Guillaud, 2020).

There are ethical issues around the question of inclusion and exclusion in analytics. Most often, these are put in the form of concerns about biased algorithms. But arguably, the question of inclusion in analytics ought to be posed more broadly. For example, Emily Ackerman (2019) reports of having been in a wheelchair and blocked from existing an intersection by a delivery robot waiting on the ramp. This isn't algorithmic bias per se but clearly the use of the robot excluded Ackerman from an equal use of the sidewalk.

New types of artificial intelligence lead to new types of interaction. In such cases, it is of particular importance to look at the impact on traditionally disadvantaged groups. “There is increasing recognition that harnessing technologies such as AI to address problems identified by working with a minority group is an important means to create mainstream innovations. Rather than considering these outcomes as incidental, we can argue that inclusive research and innovation should be the norm” (Coughlan et al., 2019a; p.88).

Above, we discussed the ethics of surveillance itself. Here, we address the wider question of the surveillance culture. This refers not only to specific technologies, but the creation of a new social reality. “Focusing on one particular identification method misconstrues the nature of the surveillance society we’re in the process of building. Ubiquitous mass surveillance is increasingly the norm” (Schneier, 2020). Whether in China, where the infrastructure is being built by the government, or the west, where it’s being built by corporations, the outcome is the same.

What we are finding with surveillance culture is the “elasticity” of analytics ethics (Hamel, 2016) as each step of surveillance stretches what we are willing to accept a bit and makes the next step more inevitable. The uses of streetlight surveillance are allowed to grow (Marx, 2020). Surveillance becomes so pervasive it becomes impossible to escape its reach. (Malik, 2019). And nowhere is this more true than in schools and learning. The goal is “to connect assessment, enrollment, gradebook, professional learning and special education data services to its flagship student information system” (Wan, 2019). Or, as Peter Greene (2019) says, “PowerSchool is working on micromanagement and data mining in order to make things easier for the bosses. Big brother just keeps getting bigger, but mostly what that does is make a world in which the people who actually do the work just look smaller and smaller.”

Audrey Watters captures the issue of surveillance culture quite well. It’s not just that we are being watched, it’s that everything we do is being turned into data for someone else’s use – often against us. She says “These products – plagiarism detection, automated essay grading, and writing assistance software – are built using algorithms that are in turn built on students’ work (and often too the writing we all stick up somewhere on the Internet). It is taken without our consent. Scholarship – both the content and the structure – is reduced to data, to a raw material used to produce a product sold back to the very institutions where scholars teach and learn.” (Watters, 2019). As Watters writes, “In her book *The Age of Surveillance Capitalism*, Shoshana Zuboff calls this ‘rendition,’ the dispossession of human thoughts, emotions, and experiences by software companies, the reduction of the complexities and richness of human life to data, and the use of this data to build algorithms that shape and predict human behavior.”

The products that depend on analytics engines – plagiarism detection, automated essay grading, and writing assistance software – are built using algorithms that are in turn built on students’ work. And this work is often taken without consent, or (as the lawsuit affirming TurnItIn’s right to use student essays) consent demanded as an educational requirement (Masnick, 2008). And “Scholarship – both the content and the structure – is reduced to data, to a raw material used to produce a product sold back to the very institutions where scholars teach and learn.” (Watters, 2019) And in a wider sense, everything is reduced to data, and the value of everything becomes the value of that data. People no longer simply create videos, they are “influencers”. Courses are no longer locations for discussion and learning, they produce “outcomes”.

There is the sense that analytics and AI cannot reason, cannot understand, and therefore cannot know the weight of their decisions. This, somehow, must be determined. But as Brown (2017) asks, “Who gets to decide what is the right or wrong behaviour for a machine? What would AI with a conscience look like?” On the other hand, perhaps AI can learn the difference between right and wrong for itself. Ambarish Mitra (2018) asks, “What if we could collect data on what each and every person thinks is the right thing to do? ... With enough inputs, we could utilize AI to analyze these massive data sets – a monumental, if not Herculean, task – and drive ourselves toward a better system of morality... We can train AI to identify good and evil, and then use it to teach us morality.” The danger in this is that people may lose the sense of right and wrong, and there are suggestions that this is already happening. Graham Brown-Martin argues, for example, “At the moment within social media platforms we are seeing the results of not having ethics, which is potentially very damaging.” (Clement-Jones et al, 2018; para247). Do right and wrong become what the machine allows it to be? This is perhaps the intuition being captured by people who are concerned that AI results in a loss of humanity. And when we depend on analytics to decide on right and wrong, what does that do to our sense of morality?

While it may be intuitive to argue that human designers and owners ought to take responsibility for the actions of an AI, arguments have been advanced suggesting that autonomous agents are responsible in their own right, thereby possibly absolving humans of blame. “Emerging AI technologies can place further distance between the result of an action and the actor who caused it, raising questions about who should be held liable and under what circumstances.” (Fjeld et al., 2020; p.34)

The argument from AI autonomy has a variety of forms. In one, advanced (tentatively) by the IEEE. It draws the distinction between “moral agents” and “moral patients” (or “moral subjects”) to suggest that we ought to distinguish between how an outcome occurred, and the consequence of that outcome, and suggests that autonomous self-organizing systems

may operate independently of the intent of the designer (IEEE, 2016; p.196) As Bostrom and Yubkowsky (2020) write, “The local, specific behavior of the AI may not be predictable apart from its safety ,even if the programmers do everything right.” It may seem unjust to hold designers responsible in such cases.

Focus on Ethical Issues

In this section we examine the ethical issues being addressed by codes of conduct. Most often these are not stated explicitly, but must be inferred from the sorts of behaviours or outcomes being expressly discussed.

The Good that Can Be Done

While ethical codes are typically thought of as identifying wrongs, in the sense of “thou shalt not”, it should be noted that many codes reference first the “good” that can be accomplished by the discipline or profession being discussed. This is especially the case in relation to data management and data research, which are new fields, and where the benefits may not be immediately obvious.

For example, while the United Kingdom Data Ethics Framework “sets out clear principles for how data should be used in the public sector,” it is with the intention to “maximise the value of data whilst also setting the highest standards for transparency and accountability when building or buying new data technology” (Gov.UK, 2018), advising researchers to “start with clear user need and public benefit.” Also in the U.K., the list of principles outlines by the House of Lords Select Committee on AI principles reflect a purpose “for the common good and benefit of humanity” including privacy rights, the right to be educated, “to flourish mentally, emotionally and economically alongside artificial intelligence” (Clement-Jones et al, 2018; para417).

Similarly, the Sorbonne Declaration (2020) points to “the benefit of society and economic development” that accrues as a result of data research. It is motivated by the good that can be done and “recognises the importance of sharing data in solving global concerns – for example, curing diseases, creating renewable energy sources, or understanding climate change” (Merett, 2020). In some cases, the emphasis is on being able to be *more* ethical. The Society of Actuaries, “AI provides many new opportunities for ethical issues in practice beyond current practices, for example, ‘black box’ decision models, masked bias, and unregulated data” (Raden, 2019; p.9), all issues that received much less attention in the days before analytics.

In the field of learning analytics, there is often an explicit linkage drawn between the use of data and benefits for students, and thereby, of helping society benefit from education generally. The Open University, for example, asserts that the purpose of collecting data

should be “to identify ways of effectively supporting students to achieve their declared study goals” (OU, 2014; p.4.2.2). The Asilomar Convention for Learning Research in Higher Education principles were based on “the promise of education to improve the human condition”, as expressed by two tenets of educational research: to “advance the science of learning for the improvement of higher education”, and to share “data, discovery, and technology among a community of researchers and educational organizations” (Stevens & Silbey, 2014).

Academic or Professional Freedom

Ethical codes frequently point to the need for freedom or autonomy for the profession. Not surprisingly, the concept of academic freedom surfaces frequently in academic codes of ethics. It is seen as something that needs to be nurtured and protected. Thus, for example, one university’s code of ethics asserts that the defence of academic freedom is an “obligation” on faculty members, stating, “it is unethical for faculty members to enter into any agreement that infringes their freedom to publish the results of research conducted within the University precincts or under University auspices... they have the obligation to defend the right of their colleagues to academic freedom. It is unethical to act so as deliberately to infringe that freedom” (SFU, 1992). Or, good practices are those that defend academic freedom (EUI, 2019).

But university professors are not alone in asserting professional independence. Researchers generally, and especially early-career researchers (ECR) “are being pressured into publishing against their ethics because of threats relating to job security” (Folan, 2020). Librarians declare that they are “explicitly committed to intellectual freedom and the freedom of access to information. We have a special obligation to ensure the free flow of information and ideas to present and future generations” (ALA, 2008). Doctors and nurses also declare the caregiver’s right to “be free to choose whom to serve, with whom to associate, and the environment in which to provide medical care” (AMA, 2001). The same assertions of independence and autonomy can be found in journalists’ code of ethics (NUJ, 2011).

Conflict of Interest

The idea that a person would use their position to personally benefit from their position of privilege or responsibility, whether directly or through the offer of gifts or benefits, is expressly prohibited by many (but by no means all) codes of ethics (CFA, 2019; IEEE, 2020; p.7.8; SFU, 1992; CPA, 2017). Different sorts of conflict of interest are mentioned by different codes of ethics.

Some codes focus on material benefits. For example, codes of ethics in the financial sector often express prohibitions against insider trading (specifically, members that “possess

material nonpublic information that could affect the value of an investment must not act or cause others to act on the information” and against “practices that distort prices or artificially inflate trading volume with the intent to mislead market participants” (CFA, 2019). Public services ethics., meanwhile, address conflict of interest as a matter of trust where the principles include “taking all possible steps to prevent and resolve any real, apparent or potential conflicts of interest,” as well as “effectively and efficiently using the public money, property and resources managed by them” (TBS 2011).

Other codes focus on integrity. We see this in professions like journalism, where “professional integrity is the cornerstone of a journalist’s credibility” (SPJ, 1996) and journalists are urged “to remain independent (and therefore avoid conflict of interest), and to be accountable” (SPJ, 2014). The primary focus of the New York Times Ethical Journalism Guidebook is avoidance of conflict of interest, and it addresses exhaustively the ways in which a journalist could be in a real or perceived conflict of interest, and counsels against them, while allowing for certain exceptions (NYT, 2018).

In education and the helping professions the codes focus on exploitation (IUPSIS, 2008; CPA, 2017; NEA, 1975; BACB, 2014; p.6; SFU, 1992; EUI, 2019; etc.). The British Columbia Teachers Federation, for example, states that “a privileged relationship exists between members and students” and stresses the importance of refraining from exploiting that relationship” (BCTF, 2020).

Harm

The prevention of harm is a theme that arises in numerous codes of ethics. Many codes trace their origins to the written principles for ethical research originating from the Nuremberg trials in 1949 that were used to convict leading Nazi medics for their atrocities during the Second World War (Kay et al., 2012). In general, research should not risk “even remote possibilities of injury, disability, or death,” nor should the harm exceed the potential benefits of the research (USHM, 2020). What counts as harm, however, varies from code to code.

Often, the nature of harm is loosely defined. Accenture’s Universal Principles for Data Ethics (Accenture, 2016; p.5) states that practitioners need to be aware of the harm the data could cause, both directly, and through the “downstream use” of data. The principles also acknowledge that data is not neutral. “There is no such thing as raw data.” The Information Technology Industry Council urges researchers to “Recognize potentials for use and misuse, the implications of such actions, and the responsibility and opportunity to take steps to avoid the reasonably predictable misuse of this technology by committing to ethics by design. (UC Berkeley, 2019)

Discrimination and human rights violations are often cited as sources of harm (IEEE, 2020; p.9.26; NEA, 1975; IFLA, 2012; NUJ, 2011; UC Berkeley, 2019; etc.). For example, the Amnesty International and Access Now “Toronto Declaration” calls on the right to redress human rights violations caused by analytics and AI. “This may include, for example, creating clear, independent, and visible processes for redress following adverse individual or societal effects,” the declaration suggests, “[and making decisions] subject to accessible and effective appeal and judicial review” (Brandom, 2018).

Several codes, by contrast, identify exemptions and cases that will not be considered harm. For example, the U.S. “Common Rule” states that research is exempt from restrictions if it is a “benign behavioral exemption”, that is, it is “brief in duration, harmless, painless, not physically invasive, not likely to have a significant adverse lasting impact on the subjects, and the investigator has no reason to think the subjects will find the interventions offensive or embarrassing” (HHS, 2018; §46.104.2.C.ii).

Quality and Standards

Ethical codes – especially professional ethical codes – also address issues related to quality and standards. Sometimes competence is defined simply as “stewardship and excellence” (TBS, 2011) or professionalism (CFA, 2019; BACB, 2014; p.6). Or a profession may seek to restrict practice to competent practitioners, for example, preventing assistance to a “noneducator in the unauthorized practice of teaching” and preventing “any entry into the profession of a person known to be unqualified in respect to character, education, or other relevant attribute” (NEA, 1975).

The code may also seek to define and reinforce exemplary behaviours such as research integrity, scientific rigor and recognition of sources. The ethical code for behavioural analysts, for example, states that researchers must not fabricate data or falsify results in their publications, must correct errors in their publications, and not omit findings that might alter interpretations of their work (BACB, 2014; p.9.0). Similarly, “The IEEE acknowledges the idea of scientific rigor in its call for creators of AI systems to define metrics, make them accessible, and measure systems” (Feljd et al., 2020; p.59). The major sources of academic misconduct are related to the misuse of intellectual property, for example, through plagiarism, piracy, misrepresentation of authorship (“personation”), and fabrication data or qualifications (EUI, 2019; BACB, 2014; p.9.0).

What are the Limits?

Finally, some ethical codes seek to address the limits of what can be done ethically. It’s not always easy to recognize these limits; it was only after years of effort that IBM announced it would see work in general facial recognition technology, for example (Krishna, 2020). Sometimes the need for limits is stated explicitly. The purpose of the U.K. Government

Data Ethics Framework, for example, to help data scientists identify the limits of what is allowed, to help practitioners consider policy when designing data science initiatives, and to identify core ethical expectations from such projects (Gov.UK, 2018).

Some discussions (e.g. Floridi et al., 2018; note5) omit consideration of the research issues (arguing “they are related specifically to the practicalities of AI development”), however they set an important ethical standard, specifically, “to create not undirected intelligence, but beneficial intelligence” (Asilomar, 2017). In other cases, specific outcomes are undesired, for example, “We should not build a society where humans are overly dependent on AI or where AI is used to control human behavior through the excessive pursuit of efficiency and convenience” (Japan, 2019; p.4). Many “individual” researchers, meanwhile, refuse to work on military or intelligence applications (Shane & Wakabayashi, 2018).

Otherwise, the limits are related to the benefits. For example, the Information and Privacy Commissioner Ontario, Canada. Data-gathering by the state should be restricted to that which is reasonably necessary to meet legitimate social objectives, and subjected to controls over its retention, subsequent use, and disclosure. (Cavoukian, 2013). Similarly, research Ethics Boards (REB) often require that the submissions for ethics approval be accompanied with statements of scientific merit and research need.

Core Values and Priorities

The previous section addressed ethical issues being addressed by codes of conduct. It was, in a sense, addressing the “purpose” of the code “qua” code of ethics, that is, it didn’t look at the social, political or economic need for codes of ethics, but rather, sought to identify the questions for which a “code of ethics” is the answer. No code of those surveyed was designed to meet all of the purposes identified, and none of the purposes identified was specifically addressed by all of the codes surveyed. We use different ethical codes to do different things.

In this section, we will focus on the “values and priorities” that can be found in the codes. These are things that might be found in the ethical “principles” described by the code, if the code is structured that way, or the things that are explicitly described as good or desirable by the code. When people state that there is a “universal” or “general” agreement on values, it is usually with respect to a subset of the items listed here that they refer. Below we have not attempted to create a tab or values mapped to codes, as some researchers (e.g. Fjeld et al., 2020) have done, but rather, to list the values with references to relevant examples where they are asserted.

Pursuit of Knowledge

The pursuit of knowledge is identified as a core value by many academic and professional codes. For example, the SFU code of ethics, addresses faculty members first as teachers, and then as scholars. “The first responsibility of university teachers is the pursuit and dissemination of knowledge and understanding through teaching and research. They must devote their energies conscientiously to develop their scholarly competence and effectiveness as teachers” (SFU, 1992).

Similarly, the National Education Association statement (NEA, 1975) is based on “recognizes the supreme importance of the pursuit of truth, devotion to excellence, and the nurture of the democratic principles.” Nor is the pursuit of knowledge limited to academics. The Society for Professional Journalists (SPJ) code of ethics, originally derived from Sigma Delta Chi’s “New Code of Ethics” in 1926 (SPJ, 2014), asserts that the primary function of journalism, according to the statements, is to inform the public and to serve the truth.

Autonomy and Individual Value

Many codes, like National Education Association code (NEA, 1975) are based on “believing in the worth and dignity of each human being. This, though, is expressed in different ways by different codes. For example, in one code, individual development is the objective, to promote “acquisition of autonomous attitudes and behavior.” (Soleil, 1923). The AI4People (Floridi et al., 2018; p.16) adopts a similar stance.

By contrast Tom Beauchamp and James Childress’s “Principles of Biomedical Ethics” contains an extended discussion of autonomy embracing the idea of “informed consent”, which requires disclosure of information, respect for decision-making, and provision of advice where requested. A similar respect for human autonomy is demanded by the High-Level Expert Group on Artificial Intelligence (AI HLEG, 2019).

Similarly, the Belmont Report begins by identifying “respect for persons”, as a core principle which “incorporates at least two basic ethical convictions: first, that individuals should be treated as autonomous agents, and second, that persons with diminished autonomy are entitled to protection.” (DHEW, 1978; p.4).

Consent

Whether or not based in the principle of autonomy or the inherent worth of people, the principle of consent is itself often cited as a fundamental value by many ethical codes (BACB, 2014; DHEW, 1978; HHS, 2018; Drachsler & Greller, 2016; etc.). However there may be variations in what counts as consent and what consent allows.

For example, the type of consent defined by the Nuremberg Code “requires that before the acceptance of an affirmative decision by the experimental subject there should be made known to him the nature, duration, and purpose of the experiment; the method and means by which it is to be conducted; all inconveniences and hazards reasonably to be expected; and the effects upon his health or person which may possibly come from his participation in the experiment” (USHM, 2020).

Several codes are more explicit about what counts as informed consent. For example, one code requires that “researchers be transparent about the research and give research subjects the choice not to participate. This includes passive data collection, such as collection of data by observing, measuring, or recording a data subject’s actions or behaviour” (IA, 2019). The same code, however, contains provisions that allow data to be collected without consent. If consent is not possible, it states, “Researchers must have legally permissible grounds to collect the data and must remove or obscure any identifying characteristics as soon as operationally possible.” There are also stipulations designed to ensure research quality and to ensure that communications about the research are accurate and not misleading (Ibid).

Meanwhile, that same code of ethics can allow the scope of consent to be extended beyond research. It is the IA Code of Standards and Ethics for Marketing Research and Data Analytics (IA, 2019). Consent is required for research purposes, but in addition “such consent can enable non-research activities to utilize research techniques for certain types of customer satisfaction, user, employee and other experience activities.” The Nuremberg Code and marketing research may stand at opposite poles of an ethical question, however, they are reflective of a society as a whole that holds consent as sacrosanct on one hand and makes legal End User Licensing Agreements (EULA) on the other hand.

Integrity

Integrity is often required of professionals (CFA, 2019; CSPL, 1995; IA, 2019; etc.), but different codes stress different aspects of integrity. The Canadian Psychological Association section on integrity speaks to accuracy, honesty, objectivity, openness, disclosure, and avoidance of conflict of interest (CPA, 2017). The European University Institute defines integrity as including such values as honesty, trust, fairness and respect. (EUI, 2019). The Ontario College of Teachers focuses on trust, which includes “fairness, openness and honesty” and integrity, which includes honesty and reliability (OCT, 2020). In Guyana, integrity includes “honest representation of one’s own credentials, fulfilment of contracts, and accountability for expenses” (Guyana, 2017). The Nolan Principles state “Holders of public office should act solely in terms of the public interest” (CSPL, 1995) while Raden (2019; p. 9) defines it as “incorruptibility”.

Confidentiality

While sometimes breaches of confidentiality are depicted as “harm”, confidentiality is often presented as a virtue in and of itself, perhaps constitutive of integrity. Thus, for example, librarians “protect each library user’s right to privacy and confidentiality with respect to information sought or received and resources consulted, borrowed, acquired or transmitted” (ALA, 2008). Similarly, the Declaration of Helsinki states that “every precaution must be taken to protect the privacy of research subjects and the confidentiality of their personal information” (WMA, 2013).

The need for confidentiality increases with the use of electronic data. The authors of a 1973 report for the U.S. Department of Health, Education and Welfare addressing the then nascent practice of electronic data management noted that “under current law, a person’s privacy is poorly protected against arbitrary or abusive record-keeping practices” (Ware et al., 1973; p.xx). Government policy, they argued, should be designed to limit intrusiveness, to maximize fairness, and to create legitimate and enforceable expectations of confidentiality (Linowes et al., 1977; pp.14-15).

Confidentiality, expressed as privacy, is a core principle for data and information services and codes regulating those. For example, the Federal Trade Commission promotes principles that “are widely accepted as essential to ensuring that the collection, use, and dissemination of personal information are conducted fairly and in a manner consistent with consumer privacy interests.” (Pitofsky et al., 1998; p.ii).

It should be noted that exceptions to confidentiality may be allowed, especially where required by law. For example, the British Columbia Teachers’ Federation code states explicitly that “It shall not be considered a breach of the Code of Ethics for a member to follow the legal requirements for reporting child protection issues” (BCTF, 2020). Similarly, in medical informatics, confidentiality can be compromised “by the legitimate, appropriate and relevant data-needs of a free, responsible and democratic society, and by the equal and competing rights of others” (IMIA, 2015).

Care

Care, which includes “compassion, acceptance, interest and insight for developing students’ potential” (OCT, 2020) is found in numerous ethical codes (CNA, 2017; CFA, 2019; IUPSIS, 2008; CPA, 2017; etc.) but is manifest differently in each code in this it appears. Contrasting the OCT definition, for example, is the Canadian Nurses Association discussion of “provision of care” references speech and body language, building relationships, learning from “near misses”, adjusting priorities and minimizing harm, safeguarding care during job actions, and more. It is worth noting that the promotion of

dignity means to “take into account their values, customs and spiritual beliefs, as well as their social and economic circumstances without judgment or bias.” (CAN, 2017; p.12)

The National Council of Educational Research and Training is almost unique in an assertion of care, in the explanatory notes, that states “the demonstration of genuine love and affection by teachers for their students is essential for learning to happen. Treating all children with love and affection irrespective of their school performance and achievement level is the core of the teaching learning process” (NCERT, 2010).

Other codes (e.g. CFA, 2019) adopt a more legalist interpretation of “duty of care”, for example, that researchers must “prioritize data subject privacy above business objectives, be honest, transparent, and straightforward in all interactions (and respect the rights and well-being of data subjects” (IA, 2019). Meanwhile there is a sense of “care” that means “diligence and rigor”; this is the sense intended in the Nuremberg Code (USHM, 2020) and the American Medical Association (Riddick, 2003).

Competence and Authority

Many of the codes identify competence or authority to practice in the profession as core values or principles (CFA, 2019; IEEE, 2020; p.7.8; IUPSYS, 2008; etc.). This is expressed in several ways: members of the profession may be expected to perform in a competent manner, or they may be required to remain within their domain of competence, or they may be obligated to ensure that unqualified people do not practice the profession (NEA, 1975, as cited above).

For example, behaviour analysts are expected to rely on scientific evidence and remain within the domain of their competence (BACB, 2014; p.6). Similarly, the Nuremberg Code also determines that the researcher should be a qualified scientist and that the research ought to have scientific merit and be based on sound theory and previous testing (USHM, 2020). And the CPA code (2017) requires that the practitioner be competent.

Sometimes what counts as competence is spelled out in the code. For example, the Royal Society data science ethics in government report (Drew, 2016) advises the use of robust data models in data research. Provisions in the Open University code similarly state that the modeling based on the data should be sound and free from bias, and that it requires “development of appropriate skills across the organisation” (OU, 2014; p.4.4).

Codes sometimes require that only authorized professionals perform the work. Accenture’s Universal Principles for Data Ethics (Accenture, 2016; p.5) states “practitioners should accurately represent their qualifications (and limits to their expertise).” This is especially the case where expertise is more difficult to establish or where the stakes are higher. The Guyana code of ethics for teachers, for example, requires “honest

representation of one's own credentials" (Guyana, 2017) while the Ontario Information and Privacy Commissioner Ontario states that "the authority to employ intrusive surveillance powers should generally be restricted to limited classes of individuals such as police officers" (Cavoukian, 2013).

Value and Benefit

While above we represented "the good that can be done" as aspirational, that is, something ethical codes seek to accomplish, in the present case we view the same principle as a limit, and specifically, as the research or practice must produce a benefit in order to be ethical.

In some cases, this benefit may be immediate and practical. For example, the Behavior Analyst Certification Board requires that practitioners provide "effective treatment" (BACB, 2014; p.6). It is arguable, as well, that "health-care professionals, especially, have an obligation to distinguish between remedies that represent the careful consensus of highly trained experts and snake oil" (Kennedy et al., 2002).

In other cases the requirements are more general (and more widely distributed). The Royal Society requires that researchers "show clear user need and public benefit" (Drew, 2016). Similarly, the Asilomar principles state that "AI technologies should benefit and empower as many people as possible" and "the economic prosperity created by AI should be shared broadly, to benefit all of humanity" (Asilomar, 2017). Fjeld (2020) finds a principle of "promotion of human values," and specifically, that "the ends to which AI is devoted and the means by which it is implemented should promote humanity's well being."

In other cases, the requirement that a benefit be shown is limited to requiring that practitioners demonstrate a purpose for their work. The Barcelona Principles (2010) for example require that researchers "specify purposes of data gathering in advance, and seek approval for any new uses," while the DELICATE principles require that universities "Decide on the purpose of learning analytics for your institution" and "E-xplain: Define the scope of data collection and usage" (Drachsler & Greller, 2016).

Non-Maleficence

The principle of non-maleficence is an adaptation of the principle of "do no harm" in the Hippocratic oath. This adaptation is necessary because harm is unavoidable in many circumstances; the surgeon must sometimes harm in order to heal, for example. Harm may occur in other professions as well; a teacher might punish, a researcher might violate privacy, a defence contractor might develop weapons.

So the principle of non-maleficence, as developed for example by Beauchamp and Childress (1992) means "avoiding anything which is unnecessarily or unjustifiably

harmful... (and) whether the level of harm is proportionate to the good it might achieve and whether there are other procedures that might achieve the same result without causing as much harm” (Ethics Centre, 2017). The principle arguably also requires consideration of what the *subject* considers to be harm because as Englehardt (1993) says, we engage one another as moral strangers who need to negotiate moral arrangements (Erlanger, 2002).

The definition of maleficence to be avoided can be variably broad. For example, the AMA (2001) addresses not only the nature and priority of patient care, but also “respect for law, respect of a patient’s rights, including confidences and privacy.” The AMA’s Declaration of Professional Responsibility also advocates “a commitment to respect human life” which includes a provision to “refrain from crimes against humanity (Riddick, 2003).

The principle of non-maleficence is found in numerous ethical codes, and not only medical ethics. For example, the Association for Computing Machinery (2018) states “an essential aim of computing professionals is to minimize negative consequences of computing, including threats to health, safety, personal security, and privacy,” including “examples of harm include unjustified physical or mental injury, unjustified destruction or disclosure of information, and unjustified damage to property, reputation, and the environment” (ACM, 2018).

Non-maleficence in research and data science include being minimally intrusive (Drew, 2016), to keep data secure (ibid; also Raden, 2019; p.9), to promote “resilience to attack and security, fall back plan and general safety, accuracy, reliability and reproducibility... including respect for privacy, quality and integrity of data, and access to data” (AI HLEG, 2019). AI systems, says Fjeld (2020) should perform as intended and be secure from compromise (also Drachslar & Greller, 2016).

Beneficence

Another of the principles defined by Beauchamp and Childress (1992), beneficence should be understood as more than non-maleficence and distinct from value and benefit. A professional demonstrates beneficence toward their client “not only by respecting their decisions and protecting them from harm, but also by making efforts to secure their well-being.” Moreover, “beneficence is understood in a stronger sense, as an obligation.” It’s intended as a combination of “do no harm” and “maximize benefits and minimize harm”, with the recognition that even the determination of what is harmful might create a risk of harm (DHEW, 1978; pp.6-7).

In a number of ethical codes, beneficence can be thought of as “the principle of acting with the best interest of the other in mind” (Aldcroft, 2012). This is more than merely the idea of doing good for someone, it is the idea that the role of the professional is to “prioritize”

the best interest of their client (BACB, 2015; AMA, 2001; CPA, 2017). The principle of beneficence is also raised with respect to AI (Floridi et al, 2018; p.16; Stevens & Silbey, 2014), however, in the precise statement of these principles it is unclear how they should be applied. For example, should “the common good” be included in the principle of beneficence? Should AI promote social justice, or merely be developed consistently with the principles of social justice?

Respect

The principle of respect is cited in numerous ethical codes (AMA, 2001; IUPSY, 2008; CPA, 2017; Dingwell et al., 2017; etc.), for example, acting toward students with respect and dignity (BCTF, 2020), “respect for people” (TBS, 2011), “mutual respect” (Folan, 2020), “respect for the composite culture of India among students” (NCERT, 2010), or “respect for the rights and dignity of learners” (Stevens & Silbey, 2014). Though sometimes paired with autonomy (DHEW, 1978; p.4, cited above) it is often presented quite differently. The Ontario College of Teachers code states that respect includes trust, fairness, social justice, freedom, and democracy (OCT, 2020).

Respect can also be thought of as promoting “human dignity and flourishing”, which AI4All summarizes as “who we can become (autonomous self-realisation); what we can do (human agency); what we can achieve (individual and societal capabilities); and how we can interact with each other and the world (societal cohesion)” (Floridi et al., 2018; p.7). The last two “commandments” of the Computer Ethics Institute’s Ten Commandments of Computer Ethics recommend computer professionals “think about the social consequences” and to “ensure consideration and respect for other humans” (CEI, 1992).

Democracy

Several ethical codes include “respect for democracy” among their values and principles; this can mean, variously, respect for the idea of rule by the people, respect for the “results” of democratic choice (as, say, found in public service ethics; TBS, 2011; pp.1.1-1.2), and respect for democratic values, such as justice and non-discrimination.

Democracy is also identified as both an input and output of ethical codes; the NEA code (1975) is based on “the nurture of the democratic principles,” while the Code of Professional Ethics for School Teachers in India states that “every child has a fundamental right to receive education of good quality,” where this education develops the individual personality, faith in democracy and social justice, cultural heritage and national consciousness (NCERT, 2010).

Justice and Fairness

Almost all the ethical codes consulted refer to justice in one form or another. Here it is listed alongside “fairness”, as ever since John Rawls’s influential “A Theory of Justice” (Revised, 1999) the two concepts have been linked in popular discourse, according to the principle “justice as fairness”.

As fairness, justice is cited frequently, for example, in academic codes, as fairness to students, including especially refraining from exploiting free academic labour, and ensuring credit is given for any academic work they may have depended on (SFU, 1992) and viewing academics “as role models (who) must follow a professional code of ethics” to ensure “students receive a fair, honest and uncompromising education” from teachers who “demonstrate integrity, impartiality and ethical behavior” (Guyana, 2017).

Even viewed as “fairness”, however, ambiguities remain. As the Belmont Report notes. The idea of justice, “in the sense of ‘fairness in distribution’ or ‘what is deserved’” can be viewed from numerous perspectives, each of which needs to be considered, specifically, “(1) to each person an equal share, (2) to each person according to individual need, (3) to each person according to individual effort, (4) to each person according to societal contribution, and (5) to each person according to merit.” The authors also note that exposing a disadvantaged group to risk is an injustice (DHEW, 1978; pp.6-7).

Fairness is also viewed as impartiality, an avoidance of bias or arbitrary ruling. In journalism, for example, “the primary value is to describe the news impartially – without fear or favour”, as stated by New York Times “patriarch” Adolph Ochs (NYT, 2018). Similarly, the High-Level Expert Group on Artificial Intelligence (AI HLEG, 2019) endorses “diversity, non-discrimination and fairness – including the avoidance of unfair bias, accessibility and universal design, and stakeholder participation.” And the European University Institute opposed acts that are arbitrary, biased or exploitative (EUI, 2019).

Justice, sometimes coined as “natural justice” (CPA, 2017; p.11), can also be depicted in terms of rights (Stevens & Silbey, 2014; Asilomar, 2017; Access Now, 2018). That is how it appears in the Asilomar declaration. The principles themselves reflect a broadly progressive social agenda, “compatible with ideals of human dignity, rights, freedoms, and cultural diversity,” recognizing the need for personal privacy, individual liberty, and also the idea that “AI technologies should benefit and empower as many people as possible” and “the economic prosperity created by AI should be shared broadly, to benefit all of humanity.”

This interpretation of justice is also expressed as an endorsement of diversity and prohibition of discrimination (Sullivan-Marx, 2020; Brandom, 2018; CPA, 2017; p.11;

BACB, 2014; etc.) based on various social, economic, cultural and other factors (this list varies from code to code). The National Union of Journalists code, for example, states explicitly that journalists should produce “no material likely to lead to hatred or discrimination on the grounds of a person’s age, gender, race, colour, creed, legal status, disability, marital status, or sexual orientation” (NUJ, 2011).

Justice, viewed from either the perspective of fairness or rights, can be expanded to include redress for current or past wrongs, or to prevent future wrongs. As early as 1973, U.S. Department of Health, Education and Welfare, on observing abuses in data collection, proposed a “Code of Fair Information Practice”. The intent of the code was to redress this imbalance and provide some leverage for individuals about whom data is being collected. The Toronto Declaration similarly calls for “clear, independent, and visible processes for redress following adverse individual or societal effects” (Brandom, 2018).

Depending on one’s perspective, the principle of justice may be listed together with, or apart from, any number of other principles, including fairness, rights, non-discrimination, and redress. That we have listed them here in one section does not presuppose that we are describing a single coherent core value or principle; rather, what we have here is a family of related and sometimes inconsistent principles that are often listed in the popular discourse as a single word, such as “justice”, as though there is some shared understanding of this.

Accountability and Explicability

The principles of accountability and explicability arise differently in computing and AI codes than it does in other ethical codes. In the case of academic and medical research, accountability is typically delegated to a process undertaken by a research ethics board (REB). Similarly, the Information and Privacy Commissioner of Ontario asserts that compliance with privacy rules and restrictions should be subject to independent scrutiny and that “the state must remain transparent and accountable for its use of intrusive powers through subsequent, timely, and independent scrutiny of their use” (Cavoukian, 2013).

In other disciplines, a range of additional processes describe practices such as predictability, auditing and review (Raden, 2019; p.9). As the U.S. Department of Health and Welfare argued, data should only be used for the purposes for which it was collected. And this information, however used, should be accurate; there needs to be a way for individuals to correct or amend a record of identifiable information about themselves, and organizations must assure the reliability of the data and prevent misuse of the data. These, write the authors, “define minimum standards of fair information practice” (Ware et al., 1973; p.xxi).

In digital technology, accountability also raises unique challenges. The AI4People code, for example, adds a fifth principle to the four described by Beauchamp and Childress (1992), “explicability, understood as incorporating both intelligibility and accountability” where we should be able to obtain “a factual, direct, and clear explanation of the decision-making process” (Floridi et al. 2018). As (Fjeld, 2020) summarizes, “mechanisms must be in place to ensure AI systems are accountable, and remedies must be in place to fix problems when they’re not.” Also, “AI systems should be designed and implemented to allow oversight.”

Finally, says Fjeld, “important decisions should remain under human review.” Or as Robbins (2019) says, “Meaningful human control is now being used to describe an ideal that all AI should achieve if it is going to operate in morally sensitive contexts.” As Robbins argues, “we must ensure that the decisions are not based on inappropriate considerations. If a predictive policing algorithm labels people as criminals and uses their skin color as an important consideration then we should not be using that algorithm.”

Openness

Many of the codes of ethics, especially those dedicated to research, express openness as a core value, though often with conditions attached. The Sorbonne Declaration, for example, states “research data should, as much as possible be shared openly and reused, without compromising national security, institutional autonomy, privacy, indigenous rights and the protection of intellectual property” (Sorbonne Declaration, 2020). Similarly, the Declaration of Helsinki states “researchers have a duty to make publicly available the results of their research on human subjects and are accountable for the completeness and accuracy of their reports” (WMA, 2013).

Another project, FAIRsFAIR, is based on the the FAIR Guiding Principles (GoFAIR, 2020) for scientific data management and stewardship (Wilkenson et al., 2016). The principles (and the acronym derived from them) are “Findability, Accessibility, Interoperability, and Reusability – that serve to guide data producers and publishers as they navigate around these obstacles, thereby helping to maximize the added-value gained by contemporary, formal scholarly digital publishing.”

In many cases, openness is described in terms of access serving the public good. The Asilomar Convention includes a principle of openness representing learning and scientific inquiry as “public goods essential for wellfunctioning democracies” (Stevens & Silbey, 2014). Citing The Research Data Alliance’s 2014 “The Data Harvest Report” the Concordat Working Group, (2016) authors write “the storing, sharing and re-use of scientific data on a massive scale will stimulate great new sources of wealth” (Genova et al., 2014).

Openness is also described in some principles as openness of access to services. The IFLA (2019), for example, expresses “support for the principles of open access, open source, and open licenses” and “provision of services free of cost to the user.” The Canadian Nurses association code includes “advocating for publicly administered health systems that ensure accessibility, universality, portability and comprehensiveness in necessary health-care services” (CAN, 2017).

Openness is also described in some principles as “transparency” of methods and processes (IA, 2019; Raden, 2019; p.9; Cavoukian, 2013; CSPL, 1995) in a way that often references accountability (as referenced above). The Accenture code, for example, urges professionals to foster transparency and accountability (Accenture, 2016; p.5). The High-Level Expert Group on Artificial Intelligence (AI HLEG) also advocates transparency, which includes traceability, explainability and communication.

Finally, openness can be thought of as the opposite of secrecy, as mentioned in the Department of Health, Education and Welfare report, stating that individuals should have a way to find out what information about them is in a record and how it is used (Ware et al., 1973). It is also the opposite of censorship (IFLA, 2019; ALA, 2008).

Common Cause / Solidarity

Many codes of ethics also explicitly endorse an advocacy role for professionals to promote the values stated in the code. The AMA Declaration of Professional Responsibility, for example, asserts a commitment to “advocate for social, economic, educational, and political changes that ameliorate suffering and contribute to human well-being” (Riddick, 2003).

The codes vary from advice to “teach what uplifts and unites people and refuse to be, in any way whatsoever, the propagandists of a partisan conception” (Soleil, 1923) to establishing a shared vision of teaching and to “to identify the values, knowledge and skills that are distinctive to the teaching profession” (OCT, 2016) to expressing solidarity with other members of the profession, for example, stating that criticism of other members will be conducted in private (BCTF, 2020).

Obligations and Duties

As Feffer (2017) observes, our duties often conflict. For example, we may read, “As a representative of the company, you have one set of responsibilities. As a concerned private citizen, you have other responsibilities. It’s nice when those converge, but that’s not always the case.”

We might think, for example, that a practitioner always has a primary duty to their client. Thus a doctor, lawyer, or other professional tends to the interests of the client first. A look at practice, however, makes it clear this is not the case. A doctor may (in some countries) refuse to perform a service if a patient cannot pay. An educator may be required to report on a student's substance abuse problem or immigration status.

And often, the locus of duty is not clear. For example, if a company is skewing the data used in order to sway a model toward a particular set of outcomes, does an employee have a duty to disclose this fact to the media? There may be some cases where a company is legally liable for the quality of its analytics, while in other cases (such as marketing and promotion) the requirement is less clear.

If we widen our consideration beyond simple transactions, the scope of our duties widens as well. Our duty to travel to Africa to support a learning program may not conflict with a duty to preserve the environment for people who have not yet been born. (Saugstad, 1994; Wilkinson & Doolabh, 2017) Or our desire to eat meat may conflict with what activists like Peter Singer might consider a duty to animals (Singer, 1979).

This section we look briefly at the different entities to which different codes argue that we owe allegiance, loyalty, or some other sort of obligation or duty.

Self

Most ethical codes abnegate serving or benefitting oneself, and where the self is concerned, it is typically in the service of the wider ethic, for example, our obligations as role models (Guyana, 2017). The Nolan principles, for example, make clear that the ethics of a member of the public service is selflessness (CSPL, 1995), though there is occasional acknowledgement of a duty to self (AMA, 2001).

And yet, many of the ethical principles described in the code could be construed as the cultivation of a better self, for example, one who is honest, trustworthy, integral, objective and open (this list varies from code to code) (IMIA, 2015; CSPL, 1995; CPA, 2017; IA, 2017; AITP, 2017; etc.) as well as “self-knowledge regarding how their own values, attitudes, experiences, and social contexts influence their actions, interpretations, choices, and recommendations” (IMIA, 2015).

And some principles might be thought of as promoting some desirable attributes of self, even if referring to these in others: autonomous self-realisation, human agency, and individual capabilities, for example (Floridi et al., 2018; p.7), or to “participate in programmes of professional growth like in-service education and training, seminars, symposia workshops, conferences, self study etc.” (Mizoram, 2020).

Less Fortunate

We included a place-holder for duties or obligations to the less fortunate because of an earlier reference to Peter Singer's (2009) "The Life You Can Save". Statements of any obligation toward the poor or less fortunate are impossible to find in any of the ethical codes, however, with the exception of references to specific clients of a profession, as discussed below).

That is not to say that the less fortunate are completely omitted from ethical codes. As far back as Hammurabi's Code is the edict, "the strong may not oppress the weak" (Gilman, 2005; p.4n3). At the same time, the resistance to considering such matters is telling, as summarized here: "Advocates have urged that considerations for the poor, illegal immigrants, rain forests, tribal rights, circumcision of women, water quality, air quality and the right to sanitary facilities be put into codes for administrators. As important as these issues might be they distort the purpose of ethics codes to the point that they are confusing and put political leadership in the position of quietly undermining them" (Ibid; p.47).

Student

Ethical codes for teachers or academics often specify obligations or duties to students, though in different ways. For example, "Le code Soleil" assigns a three-fold responsibility to teachers: to train the individual, the worker, and the citizen. Education, according to the code, "is the means to give all children, whatever their diversity, to reach their maximum potential" (Soleil, 1923). The National Education Association code urges teachers to "strive to help each student realize his or her potential as a worthy and effective member of society" (NEA, 1975). Further, the Open University code asserts that "students should be engaged as active agents in the implementation of learning analytics (e.g. informed consent, personalised learning paths, interventions" (OU, 2014; p.4.3.2).

Parent or Guardian, Children

Parents stand in two roles in codes of ethics. The first is to act as a proxy for children with respect to matters of consent (Kay et al., 2012). The second is as special interests that need to be protected; for example, an Indian code of ethics advises teachers to "refrain from doing any thing which may undermine students confidence in their parents or guardians" (NCERT, 20910; Mizoram, 2020) and with whom teachers need to maintain an open and trusting relationship (OCT, 2020).

Data collection began early in the field of digital media, with the FTC noting that "The practice is widespread and includes the collection of personal information from even very young children without any parental involvement or awareness" (Ibid; p.5) It is worth

noting that the principles are designed specifically to protect consumers, and that they are addressed specifically toward industry (Pitofsky et al., 1998; p.ii).

In the IEEE code there is a detailed section on “working with children” that contains provisions on safety and security, confidentiality, and whistle-blowing, noting specifically that “Adults have a responsibility to ensure that this unequal balance of power is not used for their personal advantage” (IEEE, 2017). Finally, “the Information Technology Industry Council has joined the conversation around children’s rights with a focus on emerging technologies, publishing a list of principles to guide the ethical development of artificial intelligence (AI) systems” (UC Berkeley, 2019).

Client

In many ethical codes the first and often only duty is to the client. This is especially the case for service professions such as finance and accounting, legal representation, where this is expressed as fiduciary duties, which are “special obligations between one party, often with power or the ability to exercise discretion that impacts on the other party, who may be vulnerable” (Wagner Sidlofsky, 2020).

In health care the needs of the client are often paramount. For example, the Declaration of Helsinki (WMA, 2013) states “The health of my patient will be my first consideration,” and cites the International Code of Medical Ethics in saying, “A physician shall act in the patient’s best interest when providing medical care.” It is thus “the duty of the physician to promote and safeguard the health, well-being and rights of patients, including those who are involved in medical research” (Ibid). In cases where multiple duties are owed, the client may be assigned priority, as in the case of medical research codes. “When research and clinical needs conflict, prioritize the welfare of the client” (BACB, 2014).

There is ambiguity in the concept of client, particularly with respect to the idea that the duty is to the client because the client is the one paying the bills. When care is paid by insurance, or through government programs, or corporate employers, the service recipient and the payer may be two distinct. Similarly, in digital media, costs may be paid by advertisers or publishers, who may then assert moral priority. (Done, 2010). However, as Luban (2018; p.187) argues, “‘who pays the whistler calls the tune’ is not a defensible moral principle.”

Research Subject

Research ethics codes commonly describe a duty of the researcher to the research subject, beginning with the Nuremberg Principles and established throughout the practice thereafter. The responsibilities to research participants include informed consent,

transparency, right to withdraw, reasonableness of incentives, avoidance and mitigation of harm arising from participation in research, and privacy (BERA, 2018).

In the field of data research and analytics this principle is often retained. Accenture's universal principles for data ethics, for example, state that the highest priority is "the person behind the data" (Accenture, 2016; p.5). Similarly, the Insights Association code (2019) states "respect the data subjects and their rights." In journalism, as well, "ethical journalism treats sources, subjects, colleagues and members of the public as human beings deserving of respect" (SPJ, 2014).

Employer or Funder

Public service employees are not surprisingly obligated to their employer. "Members of the public service... are tasked with "loyally carrying out the lawful decisions of their leaders and supporting ministers in their accountability to Parliament and Canadians" (TBS, 2011; pp.1.1-1.2)

The same sometimes holds true in the case of ethical codes for teachers. They may be required to "cooperate with the head of the institution and colleagues in and outside the institution in both curricular and co-curricular activities" and that a teacher should "recognize the management as the prime source of his sustainable development" (Mizoram, 2020) or to "abide by the rules and regulations established for the orderly conduct of the affairs of the University" (SFU, 1992).

The same may apply for employees in the private sector. Information technology professionals, for example, may be asked "to guard my employer's interests, and to advise him or her wisely and honestly" (AITP, 2017). Journalists, as well, are subject to obligations to the newspaper (NUJ., 1936). Even funders may make a claim on the duties of the researcher (Dingwell et al., 2017).

Colleagues, Union or Profession

Professional associations and unions frequently include loyalty to the professional association or union as a part of the code of ethics, either explicitly, or expressed as an obligation owed to colleagues (NUJ, 1936; AITP, 2017; SFL, 1992; NEA, 1975; etc.). This is related to the idea that members are forming a voluntary association. "If a member freely declares (or professes) herself to be part of a profession, she is voluntarily implying that she will follow these special moral codes. If the majority of members of a profession follow the standards, the profession will have a good reputation and members will generally benefit" (Weil, 2008).

Stakeholders

The term *stakeholders* is sometimes used without elaboration to indicate the presence of a general duty or obligation (BERA, 2018). Fjeld (2020) asserts for example that “developers of AI systems should make sure to consult all stakeholders in the system and plan for long-term effects.” The Open University policy is based on “significant consultation with key stakeholders and review of existing practice in other higher education institutions and detailed in the literature” (OU, 2014; p.1.2.6). Similarly, one of the DELICATE principles (Drachsler & Greller, 2016) requires researchers “talk to stakeholders and give assurances about the data distribution and use.”

What is a stakeholder? It expands on the concept of “stockholder” and is intended to represent a wider body of interests to which a company’s management ought to be obligated (SRI, 1963). Freeman (1984; p.25) defines it as “any group or individual who can affect, or is affected by, the achievement of a corporation’s... or organization’s purpose... or performance”. He bases it on “the interconnected relationships between a business and its customers, suppliers, employees, investors, communities and others who have a stake in the organization” (Leddecky, 2020). There are many definitions of “takeholder” (Miles, 2017; p.29) and no principled way to choose between them.

Publishers and Content Producers

Librarians are subject to special obligations to publishers, according to some codes. For example, “Librarians and other information workers’ interest is to provide the best possible access for library users to information and ideas in any media or format, whilst recognising that they are partners of authors, publishers and other creators of copyright protected works” (IFLA, 2012).

This responsibility is extended in other fields as a prohibition against plagiarism (EUI, 2019; BACB, 2014; SPJ, 2014; NUJ, 2011; NYT, 2017; etc.) and taking credit for the work of others (AITP, 2017; IEEE, 2020; BACB, 2014; etc.).

Society

References to a responsibility to society are scarce, but they do exist. BERA (2018) argues for a responsibility to serve the public interest, and in particular, responsibilities for publication and dissemination. The “Nolan principles”, (CSPL, 1995) state “Holders of public office are accountable to the public for their decisions and actions and must submit themselves to the scrutiny necessary to ensure this.”

In the field of data analytics, the last two of the Computer Ethics Institute “Ten Commandments” recommend computer professionals “think about the social consequences” and to “ensure consideration and respect for other humans” (CEI, 1992).

Though as Metcalf (2014) notes, “it appears to be the only computing ethics code that requires members to proactively consider the broad societal consequences of their programming activities” (my italics). Subsequently, the Royal Society (Drew, 2016) recommended data scientists “be alert to public perceptions.”

Law and Country

Although it has been established that there is not an ethical duty to obey an unethical law, a number of ethical codes nonetheless include respect for the law in one way or another, for example, in reporting child protection issues (BCTF, 2020), compliance with law as an “overarching principle” (IA, 2019), or “operate within the legal frameworks (and) refer to the essential legislation” (Drachsler & Greller, 2016).

Meanwhile, the Association of Information Technology Professionals Code of Ethics asserts “I shall uphold my nation and shall honor the chosen way of life of my fellow citizens,” though it is no longer extant and as Metcalf (2016) comments, “it is decades old and has some anachronisms that clash with globalized ethos of computing today.” Despite this, it was cited (in EDUCAUSE Review) as recently as 2017 (Woo, 2017).

Environment

The environment is rarely mentioned in ethical codes, though it appears in a statement of obligations to “society, its members, and the environment surrounding them” (ACM, 2018) and as “societal and environmental wellbeing – including sustainability and environmental friendliness, social impact, society and democracy” (AI HLEG, 2019).

Bases for Values and Principles

What grounds these codes of ethics? On what basis do their authors assert that this code of ethics, as opposed to some hypothetical alternative, is the code of ethics to follow? A typical explanation might be that “An individual’s professional obligations are derived from the profession and its code, tradition, society’s expectations, contracts, laws, and rules of ordinary morality” (Weil, 2008), but a closer examination raises as many questions as it answers.

Universality

Many codes simply assert that the principles embodied in the code are universal principles. Universality may be seen as a justification for moral and ethical principles; if the principle is believed by everyone, then arguably it should be believed here.

For example, the Universal Declaration of Ethical Principles for Psychologists asserts, “The Universal Declaration describes those ethical principles that are based on shared human values” (IUPSIS, 2008). It later asserts “Respect for the dignity of persons is the most

fundamental and universally found ethical principle across geographical and cultural boundaries, and across professional disciplines” (Ibid). So we see here universality being asserted as a foundation underlying a set of ethical principles. Similarly, the Asolomar Convention states that “Virtually all modern societies have strong traditions for protecting individuals in their interactions with large organizations... Norms of individual consent, privacy, and autonomy, for example, must be more vigilantly protected as the environments in which their holders reside are transformed by technology” (Stevens & Silbey, 2014).

Additional studies, such as Fjeld et al. (2020) that suggest that we have reached a consensus on ethics and analytics. We argue that this is far from the case. The appearance of “consensus” is misleading. For example, in the Fjeld et al., survey, though 97% of the studies cite “privacy” as a principle, consensus is much smaller if we look at it in detail (Ibid; p.21), and the same if we look at the others, e.g. Accountability (Ibid; p.28). Assertions of universality made elsewhere (for example: Pitofsky, 1998; p.7; Singer & Vinson, 2002; CPA, 2017; Raden, 2019; p.11) can be subject to similar criticisms.

In their examination of teacher codes of ethics, Maxwell and Schwimmer (2016) found “analysis did not reveal an overlapping consensus on teachers’ ethical obligations.” Nor are they alone in their findings; citing Campbell (2008; p.358) they observe that “despite extensive research on the ethical dimensions of teaching, scholars in the field do not appear to be any closer to agreement on ‘the moral essence of teacher professionalism’.” Similarly, Wilkinson (2007; p.382) “argues that the teaching profession has failed ‘to unite around any agreed set of transcendental values which it might serve’.” And van Nuland and Khandelwal (2006; p.18) report “The model used for the codes varies greatly from country to country.” The selection below is a sample; many more codes may be viewed in the EITCO website (IIEP, 2020).

Fundamental Rights

The High-Level Expert Group on Artificial Intelligence cites four ethical principles, “rooted in fundamental rights, which must be respected in order to ensure that AI systems are developed, deployed and used in a trustworthy manner” (AI HLEG, 2019).

As noted above, the Access Now report specifically adopts a human rights framework “The use of international human rights law and its well-developed standards and institutions to examine artificial intelligence systems can contribute to the conversations already happening, and provide a universal vocabulary and forums established to address power differentials” (Access Now, 2018; p.6).

The Toronto Declaration “focuses on the obligation to prevent machine learning systems from discriminating, and in some cases violating, existing human rights law. The declaration was announced as part of the RightsCon conference, an annual gathering of digital and human rights groups” (Brandom, 2018).

Nonetheless, it is not clear what these fundamental rights are. Their statement in documents such as the U.S. Bill of Rights, the Canadian Charter of Rights and Freedoms, or the Universal Declaration of Human Rights, is very different. Is the right to bear arms a fundamental right? Is the right to an education a fundamental right?

Fact

Arguments drawing from statements of fact about the world are sometimes used to support ethical principles. For example, the Universal Declaration of Ethical Principles for Psychologists asserts, “All human beings, as well as being individuals, are interdependent social beings that are born into, live in, and are a part of the history and ongoing evolution of their peoples... as such, respect for the dignity of persons includes moral consideration of and respect for the dignity of peoples” (IUPSIS, 2008).

Against such assertions of fact the “is-ought” problem may be raised. As David Hume (1739) argued, moral arguments frequently infer from what “is” the case to what “ought” to be the case, but “as this ought, or ought not, expresses some new relation or affirmation, ‘tis necessary that it should be observed and explained; and at the same time that a reason should be given” (Hume, 1888; p.469). Such “oughts” may be supported with reference to goals or requirements (see below), or with reference to institutional facts, such as laws (Searle, 1964).

Balancing Risks and Benefits

The AI4People declaration states “An ethical framework for AI must be designed to maximise these opportunities and minimise the related risks” (Floridi et al., 2018; p.7). Similarly the Concordat Working Group (2016) document is of open data with the need to manage access “in order to maintain confidentiality, protect individuals’ privacy, respect consent terms, as well as managing security or other risks.” And the AI4People starts from the premise that “an ethical framework for AI must be designed to maximise these opportunities and minimise the related risks” (Floridi et al., 2018; p.7).

The balancing of risks and benefits is a broadly consequentialist approach to ethics and therefore results in a different calculation in each application. For example, the balancing of risk and benefit found in the Common Rule is focused more specifically on biomedical research, and it has to be asked, is biomedicine the ethical baseline? “Not all research has the same risks and norms as biomedicine... there has remained a low-simmering conflict

between social scientists and IRBs. This sets the stage for debates over regulating research involving big data.” (Metcalf, 2016)

It also requires an understanding of what the consequences actually are. Four of the five principles recommended by the House of Lords Select Committee on AI represent a consequentialist approach (Clement-Jones et al, 2018; para417). But what are those consequences? The Committee quotes the Information Commissioner’s Office (ICO) as stating that there was a “need to be realistic about the public’s ability to understand in detail how the technology works”, and it would be better to focus on “the consequences of AI, rather than on the way it works”, in a way that empowers individuals to exercise their rights (Ibid; para51), but this may be unrealistic.

And perhaps ethics isn’t really a case of balancing competing interests. The Information and Privacy Commissioner in Ontario (Cavoukian, 2013) asserts that “a positive-sum approach to designing a regulatory framework governing state surveillance can avoid false dichotomies and unnecessary trade-offs, demonstrating that it is indeed possible to have both public safety and personal privacy. We can and must have both effective law enforcement and rigorous privacy protections.”

Requirements of the Profession

A requirement is a statement about what a person must believe, be or do in order to accomplish a certain objective or goal. For example, the Universal Declaration of Ethical Principles for Psychologists asserts, “competent caring for the well-being of persons and peoples involves working for their benefit and, above all, doing no harm... (it) requires the application of knowledge and skills that are appropriate for the nature of a situation as well as the social and cultural context” (IUPSIS, 2008). Similarly, the American Library Association sees its role as requiring “a special obligation to ensure the free flow of information and ideas to present and future generations” (ALA, 2008). The IFLA similarly argues that “librarianship is, in its very essence, an ethical activity embodying a value-rich approach to professional work with information” (IFLA, 2012).

The same document also later asserts that “Integrity is vital to the advancement of scientific knowledge and to the maintenance of public confidence in the discipline of psychology,” which is the same type of argument, however, the objectives are much less clearly moral principles: the “advancement of scientific knowledge” and “the maintenance of public confidence.” Such arguments often proceed through a chain of requirements; IUPSIS (2008) continues, for example, to argue that “Integrity is based on honesty, and on truthful, open and accurate communications.”

Such principles may be expressed in two ways: either derived, or conditional. The principle is derived if the antecedent is already an ethical principle. In the first IUPSYs example above, for example, “competent caring for the well-being of persons and peoples” may have been previously established as an ethical principle, from which the derived principle “working for their benefit” is also established. The principle may be expressed as a conditional that describes what is entailed on (say) joining a profession: if one is engaged in competent caring for the wellbeing of persons and peoples then this requires working for their benefit.

Against such assertions of requirements, several objections may be brought forward. The first method is to argue that the requirement does not actually follow from the antecedent; one might argue, for example that competent caring does not entail working for the person’s benefit; it may only involve following proper procedures without regard to the person’s benefit. Additionally, one might argue that the antecedent has not in fact been established; for example, one might argue that being a psychologist doesn’t involve caring at all, and might only involve addressing certain disruptions in human behaviour. A criminal psychologist might take this stance, for example.

Social Good or Social Order

Social good, however defined, may be the basis of some ethical principles. The preamble to the Society for Professional Journalists (SPJ) code of ethics states that the primary function of journalism, according to the statements, is to inform the public and to serve the truth, because “public enlightenment is the forerunner of justice and the foundation of democracy” (SPJ, 2014).

A basis in social order, however, invites relativism. People’s ethical judgements are relative (Drew, 2016). “People’s support is highly context driven. People consider acceptability on a case-by-case basis, first thinking about the overall policy goals and likely intended outcome, and then weighing up privacy and unintended consequences” (Ibid). This relativism is clear in a statement from a participant: “Better that a few innocent people are a bit cross at being stopped, than a terrorist incident – because lives are at risk.” And this relativism often reflects their own interests: “a direct personal benefit (e.g. giving personalized employment advice), benefit to a local community, or public protection” (Ibid).

“Social order” can be construed to mean national interest. We see this in ethics statements guiding public service agencies and professionals. For example, Russell T. Vought, issued a memo asserting that “Office of Management and Budget (OMB) guidance on these matters seeks to support the U.S. approach to free markets, federalism, and good regulatory practices (GRPs), which has led to a robust innovation ecosystem” (Vought, 2020). The

resulting “Principles for the Stewardship of AI Applications” included such things as public participation, public trust, and scientific integrity, but also included risk assessment and management along with benefits and costs. The document also urged a non-regulatory approach to ethics in AI. A different society might describe ethics in government very differently.

Fairness

A principle of “fairness” is frequently cited with no additional support or justification.

Often, fairness is defined as essential to the ethics of the profession. The New York Times, for example, “treats its readers as fairly and openly as possible” and also “treats news sources just as fairly and openly as it treats readers” (NYT, 2018).

Fairness may be equated with objectivity. For example, a journalist may say, “it is essential that we preserve a professional detachment, free of any whiff of bias” (NYT, 2018).

While acknowledging that “there is nothing inherently unfair in trading some measure of privacy for a benefit,” the authors of a 1973 report for the U.S. Department of Health, Education and Welfare addressing the then nascent practice of electronic data management noted that “under current law, a person’s privacy is poorly protected against arbitrary or abusive record-keeping practices” (Ware et al., 1973). Hence they proposed what they called a “Code of Fair Information Practice”.

Epistemology

The advancement of knowledge and learning is often considered to be in and of itself a moral good. For example, it is used in the Universal Declaration of Ethical Principles for Psychologists to justify the principle of integrity: “Integrity is vital to the advancement of scientific knowledge and to the maintenance of public confidence in the discipline of psychology” (IUPSYS, 2008). Epistemological justification is also found in journalistic ethics: “relationships with sources require the utmost in sound judgment and self discipline to prevent the fact or appearance of partiality” (NYT, 2018). And in the case of AI ethics, it may be simply pragmatic: “our ‘decision about who should decide’ must be informed by knowledge of how AI would act instead of us” (Floridi et al., 2018; p.21).

Against this argument, one may simply deny that knowledge and learning are moral goods, and are simply things that people do, and can often be harmful (as in “curiosity killed the cat”). More often, we see such responses couched in specific terms, asserting that seeking some particular knowledge is not inherently good, for example, knowledge related to advanced weapons research, violations of personal confidentiality, and a host of other real

or imagined harms. Seneca, for example, argued “This desire to know more than is sufficient is a sort of intemperance” (Letter 88:36).

Trust

In order to do any number of things, you need trust, or some of the components of trust. As a result, the elements of trust in themselves can be cited as justification for moral principles. For example, the Universal Declaration of Ethical Principles for Psychologists writes “Integrity is vital... to the maintenance of public confidence in the discipline of psychology” (IUPSYS, 2008). Chartered Financial Analysts seek to “promote the integrity and viability of the global capital markets for the ultimate benefit of society” (CFA, 2019).

Similar principles underlie ethics in journalism; “integrity is the cornerstone of a journalist’s credibility” (SPJ, 1996).

Similarly, the New York Times asserts, “The reputation of The Times rests upon such perceptions, and so do the professional reputations of its staff members.” If we here interpret “public confidence” as an aspect of trust, we see how the authors are appealing to the principle of trust to support the assertion that integrity is a moral principle.

Against this it may be argued that trust is neither good nor bad in and of itself, and indeed, that trust may be abused in certain cases, which could make measures that promote trust also bad. Moreover, it could be argued that trust is too fragile a foundation for moral principles, as it may be broken even without ill attempts. Further, it may be argued that trustless systems are in fact morally superior, because they do not create the possibility that trust may be breached, thus preserving the integrity of whatever it was that trust was intended to support.

Defensibility

Another way to define an ethical principle is to say that it is descriptive of “conduct that you (or your organization) would be willing to defend”. For example, the National Union of Journalist code of conduct (NUJ, 2011) offers “guidance and financial support of members who may suffer loss of work for conforming to union principles.”

“Through years of courageous struggle for better wages and working conditions its pioneers and their successors have kept these aims in mind, and have made provision in union rules not only for penalties on offenders, but for the guidance and financial support of members who may suffer loss of work for conforming to union principles” (NUJ, 1936).

Includes burden or onus – responding to U.S. Whitehouse – Guidance for Regulation of Artificial Intelligence Applications – Responding to these guidelines, the American Academy of Nursing argued for a less business focused assessment of the risks and benefits

of AI, saying “federal agencies should broaden the concept around use of AI related social goals when considering fairness and non-discrimination in healthcare.” They also urged that “federal agencies consider patient, provider, and system burden in the evaluation of AI benefits and costs” and “include data accuracy, validity, and reliability” in this assessment (Sullivan-Marx, 2020)

Results of the Study

After having studied a certain number of codes of ethics, in the light of the applications of analytics and arising ethical issues considered above, the following statements can be asserted.

1. None of the statements address all of the issues in learning analytics extant in the literature, and arguably, all of these statements, taken together, still fail to address all these issues.
2. Those issues that they address, they often fail to actually resolve. Often the principles state what should be considered, but leave open what should be the resolution of that consideration.
3. There are legal aspects to analytics, and there are ethical aspects, and there is a distinction between the two, though this distinction is not always clear.
4. Although there is convergence around some topics of interest, there is no consensus with respect to the ethics involved.
5. In fact, there are conflicts, both between the different statements of principles, and often, between the principles themselves (often described as a need to “balance” competing principles).
6. Even were there consensus, it is clear that this would be a minimal consensus, and that important areas of concern addressed in one domain might be entirely overlooked in another domain.
7. Ethical principles and their application vary from discipline to discipline, and from culture to culture.
8. There is no common shared foundation for the ethical principles described. As we will see below, these statements of principles select on an ad hoc basis from different ethical ideas and traditions.
9. Often these principles include elements of monitoring and enforcement, thus begging the question of why or for what reason an individual would adhere to the ethical principle stated.

Concluding Remarks

It is premature (if it is possible at all) to talk about “the ethics of such and such” as though we have solved ethics. There are multiple perspectives on ethics, and these are represented in the very different ethical codes from various disciplines. Approaches based in simple principles, such as an appeal to consequences, or such as in terms of rights and duties, and as such, as statements of rules or principles, fail to address the complexity of ethics especially as regards learning and analytics. The assertion of a universal nature of ethics doesn’t take into account context and particular situations, and it doesn’t take into account larger interconnected environment in which all this takes place.

Additionally, based in simple principles don’t take into account how analytics themselves work. Analytics systems are not based on rules or principles, they are statistical, using techniques such as clustering and regression. As such, their “input” is going to be complex, and they will produce unexpected consequences in a way that reflects the complexity of humans and human society.

There is an argument, with which we are sympathetic, that when we ask ethical questions, such as “what makes so-and-so think it would be appropriate to post such-and-such?” we are not looking for a single answer, but a complex of factors based on individual identity, society, circumstances and perspective. This suggests an ethics based on different objectives – not “rights” or “fairness” but rather things like a sense of compassion or on a philosophical perspective that uses a relational and context-bound approach toward morality and decision making, for example, as found in work based in conviviality or the ethics of care.

References

- Association of Information Technology Professionals (AITP). (2017, July 11). Ethics, Code of Conduct & Conflict of Interest. Wayback Machine.
<https://web.archive.org/web/20170711191254/https://www.aitp.org/?page=EthicsCodeofConduct>
- Bannan-Ritland, B. (2003). The role of design in research: The integrative learning design framework. *Educational Researcher*, 32(1), 21-24. doi:10.3102/0013189X032001021
- Boyer, A., & Bonnin, G. (2019). *Higher Education and the Revolution of Learning Analytics*. International Council for Open and Distance Education.
https://static1.squarespace.com/static/5b99664675f9eea7a3ecee82/t/5beb449703ce644d00213dc1/1542145198920/anne_la_report+cc+licence.pdf
- Brodsky, A., Shao, G., Krishnamoorthy, M., Narayanan, A., Menascé, D., & Ak, R. (2015). Analysis and Optimization in Smart Manufacturing based on a Reusable Knowledge

Base for Process Performance Models. *Proceedings of the 2015 IEEE International Conference on Big Data (Big Data)*, 1418–1427.

<https://ieeexplore.ieee.org/abstract/document/7363902>

Drew, C. (2018). Design for data ethics: Using service design approaches to operationalize ethical principles on four projects. *Philosophical Transactions of the Royal Society A. Mathematical, Physical and Engineering Sciences*, 376(2128), 20170353.

<https://doi.org/10.1098/rsta.2017.0353>

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-based Approaches to Principles for AI*. Berkman Klein Center for Internet & Society.

https://dash.harvard.edu/bitstream/handle/1/42160420/HLS%20White%20Paper%20Final_v3.pdf?sequence=1&isAllowed=y

Hicks, M., Narayan, A., Turkle, S., & Medina, E. (n.d.). How to Recognize AI Snake Oil. Promises and Perils of AI. Retrieved November 27, 2019, from

https://docs.google.com/document/d/1s_AgoeL2y_4iuedGuQNH6Fl1744twhe8Kj2qSfTqyHg/edit?fbclid=IwAR0QS_uS-QXJB8rxgni_zGm5KU0oQPa9AJPFv-NpKcBlOkKIJZ0J4uefhg0o#heading=h.ypt4v4y21eo5 Author, A. A.,

Koehler, M. J., & Mishra, P. (2009). What is technological pedagogical content knowledge? *Contemporary Issues in Technology and Teacher Education*, 9(1), 60-70.

Wilson, G., & Stacey, E. (2004). Online interaction impacts on learning: Teaching the teachers to teach online. *Australasian Journal of Educational Technology*, 20(1), 33-48. Retrieved from <http://www.ascilite.org.au/ajet/ajet20/wilson.html>