The Ethics of Learning Analytics in Australian Higher Education

DISCUSSION PAPER

PREPARED BY

Linda Corrin

Gregor Kennedy

Sarah French

Simon Buckingham Shum

Kirsty Kitto

Abelardo Pardo

Deborah West

Negin Mirriahi

Cassandra Colvin

Table of Contents

Please cite as:
Corrin, L., Kennedy, G., French, S., Buckingham Shum S., Kitto, K., Pardo, A., West, D., Mirriahi, N., & Colvin, C. (2019). *The Ethics of Learning Analytics in Australian Higher Education. A Discussion Paper.* https://melbourne-cshe.unimelb.edu.au/research/research-projects/edutech/the-ethical-use-of-learning-analytics

Executive Summary

The ever-increasing availability of data about student activities in educational environments presents many opportunities for the improvement of learning and teaching through the use of learning analytics. In applying analytics, there is an obligation that educators and institutions ensure that data and analysis techniques are used appropriately. The range of ethical considerations that educational institutions must face is complex and many institutions are still formulating their approach to ensuring ethical practice in this field.

The purpose of this discussion paper is to explore the key ethical issues regarding the use of learning analytics and to provide guidance on how these might be considered. It has been compiled by a group of learning analytics experts to promote discussion of ethical principles, approaches to learning analytics, and topics that still require further consideration, particularly in the Australian Higher Education context.

The paper begins with a brief overview of the field of learning analytics and the emergence of the discussion around the ethical use of data within this context. This is followed by an examination of seven key ethical principles, including:

- Privacy
- Data ownership and control
- Transparency
- Consent
- Anonymity
- Non-maleficence and beneficence
- Data management and security
- Access

The paper then profiles a range of existing learning analytics frameworks used internationally including the JISC Code of Practice for Learning Analytics, the Open University's (UK) Policy on Ethical Use of Student Data for Learning Analytics, and the DELICATE Checklist. To highlight the ways in which these principles and frameworks can be applied to practice, a number of learning analytics applications that are becoming more common in educational institutions are then explored. This is followed by a range of case studies of how the ethical practice of learning analytics has been addressed to date across a range of Australian and international universities. Throughout these scenarios and case studies is a discussion of the different perspectives of a range of stakeholders including students, staff, institutions and external bodies.

The discussion paper concludes with seven key considerations for educational leaders and practitioners as *things you might need to do* to ensure ethical practice within institutions. These include the need to:

- Recognise that the ethics of learning analytics is very complex;
- Develop clear principles and guidelines on data use in learning and teaching;
- Actively engage with multiple stakeholders;
- Establish transparency and trust;
- Avoid reinventing the wheel;
- Get a move on; and
- Develop processes to revisit and recast practice.

The discussion around the ethics of learning analytics in the Australian higher education context is ongoing, and this discussion paper offers a contribution to support and encourage its progress and the engagement of a range of stakeholders.

Introduction

The increasing pervasiveness of data and analytics in higher education creates many new opportunities to support and enhance learning environments for students. As a result, there is growing interest in the use of student data for an ever-expanding range of possible improvements to teaching and learning practices. However, a major issue that is often neglected in the implementation of analytics-based initiatives is the ethics of how and why data is collected and used. We are in a post-Snowden era, and recent events such as the highly publicised mining of Facebook data to influence political advertising by Cambridge Analytica, and the growing awareness of the uses and abuses of analytics and AI for decision-making, have resulted in increasing concern among the general public about the potential misuse of data. In educational contexts, similar cases are starting to emerge in which learning system vendors are using data to conduct educational experiments without the knowledge of students, parents or teachers. For example, Pearson Education's randomised control group studies of motivational messages to 9,000 college and university students drew criticism when student data was used without their knowledge or consent. Such events call attention to the importance of having well-considered policies and guidelines to ensure that learning analytics can be used in an ethical way within educational institutions.

The ethical considerations behind the use of student data and the design of analytics systems that will use this data are complex, and, to date, few institutional approaches have addressed ethical issues in all their complexity. Although there are some promising examples of work undertaken by a number of institutions internationally, most of these efforts have emphasised high-level, one-size-fits-all approaches to ethical policies and frameworks and have been primarily driven by privacy concerns. In such a complex context a more nuanced approach to the ethics of data and analytics is needed, one that accounts for the multiple levels and uses of analytics from design processes through to implementation. While the ethical use of student data for research purposes is well covered in established institutional research ethics guidelines, protocols and processes, the ethical implications of the use of data and analytics in the day-to-day practices of educational institutions are less clear. It is this focus on how learning analytics can be ethically operationalised within the teaching and learning environment that is the concern of this discussion paper.

The goal of this discussion paper is to explore the key issues relating to the ethical use of learning analytics in Australian higher education institutions. It is authored by an expert group of learning analytics researchers and practitioners and seeks to provide an overview of the current progress that has been made in relation to ethical practice in the field, as well as to highlight some of the issues that still require further consideration. The discussion paper builds on the outcomes of several national projects that profiled the implementation of learning analytics in Australian higher education institutions and identified ethics as a key issue to be addressed in order to achieve continued growth (e.g. West et al., 2016a; Colvin et al., 2016). It also draws on work already undertaken in the field on the development of ethical frameworks for learning analytics.

This discussion paper aims to provide useful information to a range of stakeholders in the higher education sector, recognising that while there is some commonality, ethical issues may also differ for different internal and external stakeholders. For *students and student union groups* it will provide an overview of the different ways that data can be used for learning analytics and the ethical considerations that can inform their provision of consent for such practices. For *teachers* it will raise awareness of the considerations that teachers need to make in using analytics to support their practice. For *learning analytics professionals* it will inform processes for design, selection and implementation of learning analytics tools and approaches to ensure ethical practice. For *senior management* (e.g. DVCAs) and associated data governance groups it will provide guidance on the issues to be considered at a strategic and operational level to inform the development of institutional policies, processes and guidelines. Despite the need to keep these different audiences in mind, we also hope to create a common language to facilitate conversations about ethics and analytics among various stakeholders so as to progress approaches to ensuring the ethical use of learning analytics in Australian higher education institutions.

The paper begins with an exploration of the field of learning analytics and examines some of the recent scholarly work on ethical issues to provide context for the discussion. This is followed by an outline of the key ethical principles that are important in the consideration of how learning analytics are designed and used in practice. A review of existing learning analytics ethical frameworks is provided along with examples of real applications of learning analytics in educational environments. This is followed by a range of case studies that illustrate how different institutions, both nationally and internationally, have been working towards ensuring ethical practice in relation to learning analytics. The paper concludes with suggestions for educational leaders and practitioners about what needs to be considered and done to ensure that ethical practice of learning analytics is acknowledged and operationalised within educational institutions.

Background

The field of learning analytics aims to provide meaningful ways of using data to support student learning within learning environments. The Society for Learning Analytics Research (SoLAR) define learning analytics as: “The practice of developing actionable insights through the collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs”. As the use of technology increases in educational environments, greater opportunities emerge to collect data about students and their activities. Researchers and practitioners have developed and appropriated many different analysis techniques and tools, from fields such as educational data mining, machine learning and psychology, to transform this data into meaningful outputs that can be used to improve student learning. These outputs can be employed at different levels ranging from cross-institutional reporting and benchmarking (macro-level), to institution-wide analytics used to inform management and resourcing (meso-level), to analytics that are delivered directly to teachers and students in the form of reports, tools or dashboards (micro-level) (Buckingham Shum, 2012).

While the field of learning analytics offers many new possibilities to improve student support, teaching practice, and learning environments, analytics must be used in a responsible way that protects students and teachers. However, traditional ethical and legal frameworks are challenged by the rapid pace of technological change in the field of learning analytics. The development of new analytics tools and approaches has been moving faster than most institutions have been able to address the ethical issues created by these developments (Roberts et al., 2016). Similarly, West et al. (2016b) argue that “the legal and regulatory context is often slower to respond than advances in analytics-related technology would demand” (p. 906).

Ethical issues are also complicated by the fact that the forms of data available for the purposes of learning analytics are constantly expanding and the contexts in which data is being used are changing. Ethical considerations vary substantially depending upon the uses and purposes of learning analytics such as whether student data is collected and analysed with the intention to improve teaching and learning, inform institutional planning and processes, or conduct research into the student experience. Further complexities arise as a result of the different stakeholders’ interests including university managers, academics and students, who each may be interested in learning analytics for different and perhaps conflicting reasons.

Not long after the field of learning analytics was established, researchers began to identify growing concerns in relation to privacy and ethics (Siemens, 2013; Slade & Prinsloo, 2013; Heath, 2014; Pardo & Siemens, 2014). They pointed to the ethical implications that had emerged as a result of institutions’ increased access to student data in digital environments, and focussed predominantly on privacy concerns, as well as ethical issues relating to the ownership of student data. However, it wasn’t until the LACE (European Learning Analytics Community Exchange) project established a series of workshops on ethics and privacy in learning analytics (EP4LA) in 2014 and 2015 that ethical and privacy concerns and solutions began to emerge as key considerations (Ferguson et al., 2016). Following the 5th annual conference on Learning Analytics and Knowledge in 2015 (LAK15), which included an EP4LA workshop, there was a notable increase in scholarly literature focussed on ethical concerns, including a special issue of the Journal of Learning Analytics dedicated to ethics and privacy (Ferguson et al., 2016), and a special issue of Education Technology Research and Development exploring the relationship of ethics and privacy in learning analytics to the field of educational technology (Ifenthaler & Tracey, 2016). This literature highlights the limited attention given to ethical and privacy concerns in the design of learning analytics tools and infrastructure, emphasising the need for learning analytics applications to be designed and implemented with ethical considerations in mind (Drachsler & Greller, 2016; Gursoy et al., 2017).

In addition to the growing body of literature highlighting ethical concerns in the field of learning analytics, a number of codes of practice have been created.These codes of practice illustrate an awareness from the learning analytics community that institutions require clear guidelines to be able to progress the applications of learning analytics within legal and ethical frameworks. The Open University in the UK was one of the first universities to establish a code of practice in 2014, followed soon after by a code of practice developed by JISC in 2015, an organisation that provides advice on digital technology for education and research in the UK. These codes of practice and frameworks will be examined in more detail below.

While the codes of practice offer a useful starting point for consideration of how guidelines for the ethical use of learning analytics can be developed, Drachsler et al. (2016) note that a code of practice is “only the first step in developing a new ecosystem for learning analytics”, and argue that an architecture for large scale implementation will need to go beyond policy guidelines (p. 492). In a recent article, Lang et al. (2018) also advocate for a personal code of ethics to be employed, in addition to institutional codes, to better address the ethical responsibilities of individual practitioners of learning analytics.

A range of scholars in the field have offered sets of principles and guidelines designed to inform a framework for a ethical approach to learning analytics. In an early article Slade and Prinsloo (2013) identified six principles to help guide higher education institutions to address ethical issues in learning analytics. These principles call for the recognition of the complexity of student identity, performance and success, while making use of data in transparent and moral ways. More recently, Drachsler and Greller (2016) created the 8-point ‘DELICATE’ checklist of issues to be considered in the implementation of learning analytics, and Steiner et al. (2016) proposed the LEA’s BOX privacy and data protection framework, a project funded by the European Commission with a set of eight foundational requirements for the ethical treatment of data in learning analytics platforms.

A central focus within recent scholarship on ethics and learning analytics concerns the need to engage more directly with students (Ifenthaler & Schumacher, 2016; Arnold & Sclater, 2017) and imbue students with greater agency over their data (Prinsloo & Slade, 2016). In a series of articles, Prinsloo and Slade examine the complex issue of student consent in digital learning environments (2015) and the notion of student vulnerability (2016), and explore ways to empower students as participants in learning analytics. It has also been observed that students who have had some experience with the outputs of learning analytics systems are more comfortable with allowing their data to be used for learning analytics initiatives (Arnold & Sclater, 2017). However, studies that have directly involved students in discussions around the ethical use of their data have found that these discussions are sometimes complicated by a lack of shared language around data and analytics terminology (Brooker et al., 2017). These studies highlight the need for learning analytics designers and educational institutions to think carefully about how to frame discussions with students to ensure a shared understanding.

In an Australian education context, while learning analytics has received significant attention since the early 2010s, there has been limited consideration of ethical issues. In 2014, Heath published an article on the potential contribution of privacy theories to learning analytics and this work underpinned the development of one of the first ethical guidelines for the implementation of learning analytics in an Australian university at the University of Wollongong. The need for a greater focus on ethics in the Australian context was an outcome of two national projects which focused on the implementation of learning analytics in Australian higher education (Colvin et al., 2016; West et al., 2016). The Colvin et al. (2016) report notes that “the relative silence afforded to ethics ... is highly significant and does not reflect the seriousness with which the sector should consider these issues” (p. 33). The report recommends that a “national conversation be initiated in which ethical considerations will be identified, framed and possible actions identified” (p. 33). Drawing on the outcomes of the West et al. (2016a) project, a framework for ethical decision making was proposed to provide a model for ensuring ethical practice (West et al., 2016b).

In the literature concerning the ethics of learning analytics the distinction between ethical and legal principles is sometimes blurred. However, in developing a taxonomy of 86 ethical, legal and logistical issues for the field of learning analytics for the JISC code of practice, Sclater (2016) found that most issues comprise an overlap of both ethical and legal dimensions since “laws are often underpinned by ethics” (p. 22). Indeed, many of the ethical principles emphasised in the literature, including privacy and informed consent, are simultaneously legal and ethical concepts. On the other hand, there are also instances in which ethical and legal issues are distinct and must be addressed separately. For example, there may be times when institutions would have ethical obligations that conflict with or override legal and privacy considerations, such as if student data or analytics revealed that a student was at risk of physical harm.

As has been observed, there is a relatively low level of legal ‘maturity’ in the field of learning analytics (Colvin et al., 2015). Established legal frameworks, such as the *Privacy Act 1988* (Privacy Act), were designed prior to technological advances and are not always easily applied to the digital context. The Privacy Act regulates how personal information is handled but is only applicable to identifiable individuals. Where data is de-identified the privacy principles do not apply. However due to the ability to re-identify data in different contexts, which is possible in many educational environments, the Office of the Australian Information Commissioner advises that a risk management approach, which acknowledges that the data may become re-identifiable, be taken (OAIC, 2018). As a result of the limitations of the current laws and regulations pertaining to learning analytics, ethical considerations need to be addressed that extend beyond legal obligations and compliance issues. Ethical frameworks also need to be agile enough to anticipate and respond to unknown outcomes of data analysis.

These ethical frameworks and considerations also need to be able to cater to the different stakeholders involved in learning analytics. Some of these stakeholders have a direct involvement with the learning analytics endeavours of a university such as students, academics, analysts and institutional decision makers, while others are involved more indirectly such as vendors and government. For each of these groups, different ethical tensions may exist, and often design decisions about learning analytics are made in environments with inherent power imbalances between stakeholders (Prinsloo & Slade, 2016; Rubel & Jones, 2016). For example, vulnerable students, such as those from a low SES background, might be stereotyped as a result of algorithms and learning analytics applications. Similarly, students may not be given the option to “opt out” of the collection and analysis of their data, regardless of whether they receive any direct benefit from its use.

Academics may use analytics to help monitor and improve the learning environment for students, but at the same time have concerns about the use of data and analytics by line-managers and the institution to evaluate their teaching or course performance (West et al., 2016a). Analytics can often be used as a blunt tool for such evaluative purposes, especially if the pedagogical intent of the design of the learning environment is not taken into consideration when determining if patterns of student engagement are satisfactory (e.g. a low number of "clicks" doesn't always indicate that students weren't engaged in learning). Concern has also been raised about whether institutions are using the right data and techniques to evaluate teaching quality (Kitto, Williams & Alderman, 2019), which is further complicated by the fact that student satisfaction is assumed to correlate to learning performance despite this being found not always to be the case (Rienties & Toetenel, 2016).

While data can provide insight into teaching practices for evaluation purposes, ethical tensions emerge when the same data is used by academics for a secondary purpose such as to demonstrate excellence for promotions or awards, or for presentation or research publication. Normally, the collection of data for research purposes is subject to scrutiny through an institution's research ethics process, but the ease of access to data afforded through learning analytics systems can lead to its "repurposing" for research without being subject to the same scrutiny (Rodríguez-Triana, Martínez-Monés & Villagrá-Sobrino, 2016; Willis, Slade & Prinsloo, 2016).

One of the potentially useful purposes to which learning analytics can be put from an institutional perspective is in responding to the requirements of governments, accreditation bodies and funding agencies. Senior leaders within institutions may look to learning analytics to provide insight into ways to improve outcomes for the greatest number of students, to increase retention and to improve overall teaching and learning indicators. However, the challenge is that the benefit to the individual student can be overlooked in the service of other macro institutional concerns and reporting. Institutions may focus on the 'public' good where the 'public' is seen as the broader student cohort and the safeguarding and improvement of the institution. Alternatively, Prinsloo and Slade (2016) suggest that institutions have a fiduciary duty to use data to improve individual student success. Institutions, therefore, must balance this requirement to support student learning with their reporting requirements to government and other agencies.

Ethical questions can also be raised about the role of vendors of educational technology and learning analytics systems in relation to the interests of students, academics and institutions. Driven by a need to enhance their product, vendors often work in partnership with institutions and have access to a range of data across many institutions. Although subject to legislation and contracts which define and outline their responsibilities, there remains the potential for a range of ethical issues to arise. For example, the ownership of data collected about students' activities through the use of educational technologies hosted by vendors in the cloud can be unclear. This issue, in particular, will be explored in more depth in a later section of this discussion paper.

This overview clearly reveals the complexity of the terrain when it comes to learning analytics, big data and institutional responses. It also shows how, after a slow start, momentum has now been created by both individuals and institutions in dealing with the ethics of learning analytics in policy and practice. This discussion paper aims to further explore and provide advice on key issues on the use of data and learning analytics in the Australian higher education context. The next section starts this exploration with a consideration of key ethical principles.

Ethical Principles

A range of ethical issues emerge when thinking about the design and use of learning analytics. This section briefly defines and discusses eight prominent ethical principles that emerge in the learning analytics literature: (1) privacy; (2) data ownership and control; (3) transparency; (4) consent; (5) anonymity; (6) non-maleficence and beneficence; (7) data management and security; and (8) access.

Privacy

Ethics and privacy are often conflated in the literature on learning analytics, yet there are some important distinctions between the two terms. While ethics is a branch of philosophy that seeks to resolve moral questions around what is wrong and right, privacy relates to the right to freedom from surveillance or unauthorised disclosure of one's personal information. In other words, ethical principles are contingent and related to social and cultural conventions, whereas privacy is a legal principle and basic human right. As Drachsler and Greller (2016) explain, "ethics is a moral code of norms and conventions that exists in society externally to a person, whereas privacy is an intrinsic part of a person's identity and integrity" (p. 90). Privacy, therefore, involves self-determination in that individuals are imbued with the capacity to determine their level of privacy or disclosure of personal information. However, technological developments have increasingly impacted upon privacy and on the ways in which we understand privacy, as users are not always fully aware of the data that is being collected about them or the ways in which it is being used.

Data Ownership and Control

Data ownership refers to the possession of, control of, and responsibility for information. Questions surrounding the ownership of data include considerations of who determines what data is collected, who has the right to claim possession over that data, who decides how any analytics applied to the data are created, used and shared, and who is responsible for the effective use of data. Ownership of data also relates to the outsourcing and transfer of data to third parties. A number of scholars point to the lack of legal clarity with respect to data ownership (Jones et al., 2014; Siemens, 2013; Greller & Draschler, 2012), noting that this issue is not unique to learning analytics, but rather emerges in relation to many forms of digital interaction. With the absence of legal systems in place to address this issue, the default position has been that the "data belongs to the owner of the data collection tool [who is], typically also the data client and beneficiary" (Greller & Draschler, 2012, p.50).

Transparency

Transparency involves the provision of information about data, intentions and processes used by an organisation that a data subject can employ to inform decision making (Turilli & Floridi, 2009). In the context of learning analytics this could relate to transparency of the data collected as well as the techniques used to analyse it and generate outcomes. A lack of transparency can cause unease and distrust for data subjects. To alleviate such concerns, it is important that higher education institutions are clear and transparent about what data is collected, the purposes of the data collection and how the data will be used, as well as how the data are processed, stored and shared. It is only by being open and transparent that it becomes possible for institutions to obtain informed consent. As student-facing analytics tools become more common, additional issues arise around whether students can interpret what is being fed back to them, and whether they can act on that information in productive ways. Institutions need to develop understanding and policies around how information about data and its analysis can be presented to students in ways that are understandable to reduce the likelihood of such processes being perceived as occurring in a “black box”.

Consent

Consent involves making a contract with the data subjects by seeking their consent for their data to be collected and analysed. For consent to be valid it must be informed; therefore, the subjects should be provided with clear and transparent information on the purposes for data collection so that they are in a position to give informed consent. They should also potentially be provided with the option to opt-out of having their data collected at any stage. However, as Sclater (2016) notes, consent is not necessarily a straightforward issue since it is not always a legal requirement that consent is obtained, such as in cases where data collection is deemed to be necessary for the ‘legitimate interests’ of an organisation (p. 32). Thus, for example, the JISC code of practice notes that there may be circumstances in which students would not be permitted to opt out of having their data collected.

The notion of informed consent is one of the more challenging ethical considerations in the context of learning analytics, and one that has received increased attention in recent scholarship. West et al. (2016b) write of the uneasy relationship between ‘consent’ and ‘informed consent’ noting that these concepts are often conflated in higher education digital environments, given that students are frequently asked to sign permission for their data to be collected, but the purposes for which the data will be used may not be made explicit (p. 914). Further, as Cormack (2016) suggests, it is not always clear prior to the collection and analysis of data what correlations will emerge or what the impact on individuals will be, which makes it difficult for institutions to provide clear and transparent information for the purposes of obtaining informed consent.

Anonymity

Anonymity involves offering individuals the choice to conceal or reveal their identity and any identifying information about themselves. In the area of learning analytics anonymity may involve the de-identification of individuals prior to data sharing or analysis. However, although it is generally agreed that institutions should make every effort to anonymise data, experts have also argued that anonymity cannot always be one hundred percent guaranteed. For example, Sclater (2016) questions whether “in the age of Big Data it is ever possible to anonymise an individual’s data such that they cannot be re-identified at some stage” (p. 33). Similarly, Drachsler and Greller (2016) suggest that “anonymised data can rather easily be de-anonymised when they are merged with other information sources” (p. 94). Granting anonymity also limits the potential uses of learning analytics as it hampers or prevents meaningful bilateral communication and limits the capacity for student intervention, feedback and support.

Non-maleficence and beneficence

Non-maleficence means that no harm should be done to participants in the process of implementing and employing learning analytics. There are a range of potentially adverse effects of learning analytics that need to be understood and minimised, including the possibility of prejudice against students when they are categorised, the potential for detrimental effects to emerge from the analysis of student progress, and the discrimination that might occur in relation to interventions (Sclater, 2016). Since it is not generally possible for analytics to provide the analyst or teacher with a full understanding of an individual's learning experiences or personal circumstances, care needs to be taken in both interpreting and acting upon the outcomes of the data analysis to minimise any potential bias or negative effects.

Beneficence implies that in addition to doing no harm, learning analytics should be employed to pursue positive outcomes and benefits. In their promotion of a shared ethical literacy for learning analytics, Slade and Prinsloo (2013) urge teachers and researchers to ask, "who benefits from learning analytics and under what conditions?" (p.14). They identify students as major stakeholders, along with others within the university community responsible for delivering and supporting teaching and learning. However, as discussed earlier, stakeholders may have conflicting interests and needs when it comes to learning analytics, which complicates the notion of beneficence.

In the area of non-maleficence and beneficence, questions arise about the level of responsibility borne by institutions and individual academic staff. In particular, does the fact that institutions and teachers now have access to greater amounts of analytics-based student data and analyses make them more accountable? Recent articles have questioned whether institutions are obliged to act on their analyses of student data (Prinsloo & Slade, 2017; Griffiths et al., 2016), including their obligation to intervene when students are at risk of failing or when they are perceived to be at risk of physical and emotional harm. For example, the recent LACE review of current issues in relation to student privacy asks: 'Does the new knowledge gained bring with it a responsibility to act upon it? What are the ramifications of action or inaction?' (Griffiths et al., 2016, p. 8).

Data Management and Security

Protecting and managing data collected for the purpose of learning analytics is a key consideration for institutions. Data management relates to the storage of data and includes considerations of where data should be stored, how long it should be stored and who has access rights to the data. Employing measures to protect and safeguard access to data is also an important part of the security of data. Management and security of learning analytics data is often complicated by the fact that the data can be generated across a range of online systems which may or may not be owned or hosted by the institution. The development and promotion of policies to encourage as much consistency as possible in data management across these different platforms is advised (Farah, Vozniuk, Rodriguez Triana & Gillet, 2017).

Access

The issue of access relates to the extent to which student data should be controlled, protected and restricted. It raises questions surrounding who should have access to student data, including teaching staff and potential employers, as well as the extent to which students should be able to have access to the data held about them and the analytics performed on their data. Sclater (2016) suggests that while few institutions are in a position to provide students with all of the data held about them, in Europe students have a legal right to request information, including the data held about them and the analytics performed on their data. Institutions are therefore obliged to facilitate access if it is requested by the student as well as to give students the opportunity to correct inaccurate personal data held about them.

Review of existing learning analytics ethical frameworks

JISC - Code of Practice for Learning Analytics

In 2015 in the UK, the Joint Information Systems Committee (JISC) released a *Code of practice for learning analytics* (Sclater & Bailey, 2015) which outlines recommendations for institutions to enhance existing information management processes to ensure the responsible, effective, legal, and ethical use of learning analytics. The code begins with the recommendation that institutions should allocate legal and ethical responsibility for the collection, anonymisation, analysis, and stewardship of student data as well as the interventions that may result from learning analytics. It recommends that key stakeholders, including students, should be consulted as part of the development of implementation processes around learning analytics. To ensure the transparency of this process, the code recommends that institutions scope out the necessary data for learning analytics with reference to a set of clearly defined and communicated objectives. Related to this call for transparency is the suggestion that student consent should be sought for any interventions that result from the analytics and that students should have the ability to amend their consent at any time.

The code addresses privacy concerns by recommending that only those people within the institution with a legitimate reason to view the data should be able to access it and that care should be taken to de-identify the data where appropriate. It also suggests that privacy restrictions may need to be overridden if it is identified that a student is at risk and there may be a legal obligation for the institution to intervene. The code highlights the importance of monitoring the validity and quality of data and analytics processes carried out on student data as well as ensuring that these processes are useful and appropriate. It is recommended that students should be able to access any learning analytics, metrics and labels attributed to them through analysis of their data and that this information should be provided in meaningful and accessible formats.

The Open University's Policy on Ethical Use of Student Data for Learning Analytics

In 2014 the Open University (OU) in the UK publicly released their policy on the ethical use of student data for learning analytics (The Open University UK, 2014). This policy defines learning analytics as "the use of raw and analysed student data to proactively identify interventions which aim to support students in achieving their study goals" (p. 1). While the policy notes that all data about students' activities could potentially be used for learning analytics, it provides the constraint that this should only happen when "there is likely to be an expected benefit (which will be evaluated) to students' learning" (p. 1). The policy categorises data into (1) student characteristic data, and (2) study behaviour data, and clearly defines what data are within and outside the scope of use for the purposes of learning analytics.

The remainder of the policy is structured around eight core principles within which three key features (transparency, responsibility, effectiveness) are considered. The eight core principles include:

1. “Learning analytics is an ethical practice that should align with core organisational principles, such as open entry to undergraduate level study.
2. The OU has a responsibility to all stakeholders to use and extract meaning from student data for the benefit of students were feasible.
3. Students should not be wholly defined by the visible data or our interpretation of that data.
4. The purpose and the boundaries regarding the use of learning analytics should be well defined and visible.
5. The University is transparent regarding data collection, and will provide students with the opportunity to update their own data and consent agreements at regular intervals.
6. Students should be engaged as active agents in the implementation of learning analytics (e.g. informed consent, personalised learning paths, interventions).
7. Modelling and interventions based on analysis of data should be sound and free from bias.
8. Adoption of learning analytics within the OU requires broad acceptance of the values and benefits (organisational culture) and the development of appropriate skills across the organisation.” (The Open University UK, 2014, p.6).

The DELICATE Checklist

The DELICATE Checklist was designed to provide researchers, institutional managers and policy makers with guidance about the process of facilitating a trusted approach to implementing learning analytics. Drawing on the outcomes of a series of international ethics and privacy workshops, and with reference to several international legal frameworks (e.g. European Data Protection Directive) and codes of practice relating to research ethics (e.g. Nuremberg Code, Helsinki Declaration), the checklist aims to stimulate thinking about a broad range of ethical and privacy considerations. In particular, it was designed to help learning analytics practitioners alleviate common fears about the use of student data including: discrimination against students; violation of privacy; data ownership; lack of transparency of learning analytics systems; asymmetrical power relationships between the data controller and data subject; pressure on students to respond to automated and sometimes artificial indicators, and the re-use of data for purposes not originally intended (Draschler & Greller, 2016).

The key dimensions of the DELICATE framework are provided below and more information is available at: http://www.laceproject.eu/ethics-privacy/

“**D** etermination: Decide on the purpose of learning analytics for your institution.

E xplain: Define the scope of data collection and usage.

L egitimate: Explain how you operate within the legal frameworks, refer to the essential legislation.

I nvolve: Talk to stakeholders and give assurances about the data distribution and use.

C onsent: Seek consent through clear consent questions.

A nonymise: De-identify individuals as much as possible

T echnical aspects: Monitor who has access to data, especially in areas with high staff turn-over.

E xternal partners: Make sure externals provide highest data security standards.” (Draschler & Greller, 2016, p. 96)

Common Learning Analytics Applications and Infrastructure

In this section we introduce some of the ways in which Learning Analytics has been used by institutions to date.

Learning Analytics Dashboards

As a greater number of learning activities are delivered online, it can be hard for an educator to see what students are doing, or for students to be aware of their peers' activities. Many different commercial educational technology products now provide some form of 'dashboard' – a combination of data visualisations designed to give insight into student progress. Increasingly we see LMS vendors releasing dashboards as a part of their product (sometimes at an extra premium). Many solutions are also starting to enable the extraction of at least some data for independent use in data analysis and in-house learning analytics systems. Dashboards might be student-facing, teacher-facing, or used primarily for institutional-level analytics.

An educator might value well designed dashboard indications of student engagement that show:

- Who has yet to log in?
- Who has yet to view the resources?
- Who has contributed the most or the least to the discussion forum?
- Who has not yet written their blog post? Who has made the required number of comments on peers' blogs?
- What proportion of students downloaded the assignment at least two weeks before the deadline?
- Which videos, and sections of a video, are watched most?
- Did students who failed to complete preparation for a flipped classroom still achieve good grades?

While more advanced analytics attempt to evaluate the *quality* of the activity, dashboards are typically restricted to quantitative logs of access and time on task.

Student-facing dashboards do introduce new ethical considerations. Not surprisingly, the first products to enter the market logged and visualised the most easily identifiable aspects of student activity. It is technically easy to log student views of web pages and resources, enabling some rudimentary indicators of student engagement. These can be fed back to educators, but also to students as a way to help them reflect on their own management of learning, and to gain an awareness of what their peers are doing. When they are poorly designed, dashboards can be confusing for students, and provide them with charts of information that they might find very difficult to act on (Corrin & de Barba, 2014). What, for example, are students expected to do with a timeline of their logins, a summary of their page views, or lists of PDFs they have downloaded? There is a risk that simply providing students with relatively raw data about their activity and behaviour could be more than irrelevant to them, it could prove stressful, and this type of "feedback" could also be demotivating for students. Thus, the design and implementation of student-facing dashboards reflect an ethical software design challenge.

When designed well, a dashboard should use concepts that the student can connect to their learning. Typically this requires more advanced analytics, and integrating the dashboard with the curriculum and assessment design (i.e. the "Learning Design") can provide actionable indicators of different elements of learning (Corrin, 2019). As indicated above, a number of vendors now have student- and staff-facing dashboards as part of their product offering. For example, Blackboard has an extensive analytics suite that includes student-facing dashboards, and some preliminary studies suggest that low GPA students find them helpful (Teasley & Whitmer, 2018). In the UK, JISC has devoted substantial efforts to designing student-facing dashboards as part of their national learning analytics platform (see https://www.jisc.ac.uk/rd/projects/effective-learning-analytics), and Apereo has an OpenDashboard that can be built according to inhouse requirements.

Dashboards that are offered by commercial vendors obviously come at a cost, and as yet there is little evidence to suggest that they actually result in better learning outcomes for students. A naive dashboard implementation can carry little benefit (Corrin & de Barba, 2014), and students may not engage with the dashboard if they do not see value to their learning (e.g. constantly reminding struggling students how their progress compares to the cohort may not be very motivating)(Park & Jo, 2015). Caution has been recommended around the implementation of "one size fits all" dashboards (Roberts, Howell & Seaman, 2017; Teasley, 2017). The evidence shows that few dashboards are grounded in educational or learning sciences research into what kinds of feedback benefit different kinds of students (Jivet, et al. 2018). When designing and implementing dashboards it is important to consider the audience of the dashboard, how the learning design is integrated into the design, what prompts the user can be given to consult the dashboard, what data will be incorporated and the frequency of updates to the data. Recent literature has advocated for a more participatory approach to the design of dashboards, including the input of end users (Roberts et al., 2016), and encouraged a greater focus on evaluation of dashboard design and impact (Bodily et al., 2018).

The ways in which teachers and students interpret the information presented in dashboards can create situations where ethical issues may arise. For example, the identification of someone as "at risk" of failing a subject in a student-facing dashboard can lead to them deciding to withdraw rather than motivating them to work harder. Similarly, if a student is identified as being at risk, what duty of care does the institution have to provide support to the student? There are also issues around the accuracy of the data and analysis applied to it which can result in different interpretations of dashboard information. Students may trust the analytics as authoritative, rather than questioning them, if they are presented in an institutionally-endorsed system. All of these examples highlight the need for institutions to engage staff and students in discussions around data and information literacy to enable them to critically reflect on the information presented to them and the actions they will take in response.

Predictive Analytics

Predictive analytics refer to the use of simple and sophisticated statistical analyses of past student attributes and behaviours to make predictions about current student learning and study processes and outcomes. A predictive statistical model can be developed based on a mixture of demographic attributes and the engagement of students who participated in a previous iteration of the course, in combination with their final marks or grades. This statistical model is then used in combination with specific observations of individual students currently enrolled in the course to return some sort of prediction, which is usually a value (for example, the estimated final grade), a probability (for example, of failing the course), or simply a prediction of an ascribed category that students might fall into (e.g. “at risk”).

The typical process to create a predictive model requires some preliminary data that contain both the factors observed (e.g. student engagement, demographics, previous scores, etc.) and the factor that needs to be predicted (e.g. final mark in the course). From these elements, a wide variety of statistical methods can be used to create the final set of instructions to predict the desired value: logistic regression, linear regression, support vector machines, etc. Given this variety, institutions may use a wide range of tools from business intelligence applications (e.g. IBM Cognos, Microsoft Power BI, Tableau) to powerful statistical packages (e.g. SPSS, SAS, R or Pandas). While these tools are useful, the most effective prediction process requires human intervention independent of the tools used. That is, an expert is required to select the appropriate data to be used and to create the steps to distil them into the final predictive model. The tools mentioned above are able to automate the core computations needed to produce the models, but significant expert intervention is required in the areas of data analysis, data mining, or business intelligence.

The benefit of predictive modelling is that an outcome of a process can be predicted before it occurs. More precisely, when appropriately deployed, predictive models may provide experts and support staff with information that can be used to anticipate situations, and as result, they can be used as the basis for educational intervention or remediation. One of the most common uses of predictive modelling is to address student attrition. A prediction as to whether a student is about to abandon a course, or the institution offers the possibility of deploying support actions that may help to prevent this.

The cost of starting the development and roll out of predictive modelling is not trivial. This is, in part due to the wide variety of data analysis and modelling processes that can be used, and the need for often complex data sets to create predictive models. Moreover, it cannot necessarily be assumed that once a predictive model is created in one context, it can be automatically applied or will hold in another context (Gašević, Dawson, Rogers & Gasevic, 2016; Kuzilek, Hlosta, Herrmannova, Zdrahal & Wolff, 2015). The difficulty of establishing reproducible and generalisable predictive models is starting to receive more attention (Swenson & Duin, 2018). Thus, proper testing and redesign should be accounted for. Aside from costs related to developing and building predictive models, there is an additional cost for these models to be adopted across institutions.

There are a range of ethical issues that need to be considered when employing predictive modelling. A fundamental ethical consideration relates to the data required to produce the model; which data sources are adequate and ethically acceptable to produce the predictive model? For example, should the model include personal information about students such as health records, home addresses, the academic level of relatives, etc.? In principle, the more factors that are taken into account, the better the accuracy of the model, but clearly ethical decisions need to be made about this.

A second issue is the *accuracy* of the predictive model in its representation of the real world (not its representation of the data underpinning the model). Even if a model is statistically robust – it “fits” the data – the model remains a simplification of the real situation that it is trying to represent. As such, caution is always needed when interpreting predictive models as there is a chance that the model does not adequately represent the “real world” or indeed a real student. Faced with this ambiguity a series of ethical questions arise, including: should students have the right to see the predictions that are generated by the model? and, should an institution be held liable if a student abandons a course and when the institution has a model that predicts this is likely to occur?

Two broader concerns are emerging in other sectors beyond education. The first relates to the use of predictive models that are, necessarily, based on the *past*. The foundation for building and validating a model is historical data (for statistical models), and specifically, training data if a machine learning approach is used. The concepts of data ethics and algorithmic accountability now make very clear how systematic biases in historical data can be perpetuated in software (Australian Human Rights, 2018; Data & Society Institute, 2018). This is no less the case in education. Moreover, a progressive educational institution will be encouraging its educators to be continuously reviewing and improving curriculum, pedagogy, learning design and assessment. If a model's performance depends on future teaching and learning looking like past practice, the model's assumptions may no longer hold if these variables are changing significantly.

The second concern relates to being automatically classified – or misclassified – as a certain sort of person, and the consequences that follow from this (related to the issue of "accuracy" above). There is a growing body of examples relating to the problems of algorithmic misclassification in sectors including retail, health, finance, and criminal justice. In education, a model may classify someone, with some degree of confidence, as "likely to struggle with this course", "at risk of dropping out", etc. While, arguably, educators have always classified students in this way informally, the potential consequences of being "pigeon-holed" in this way formally and by a machine, are quite different. And when students are informed of this classification there are real concerns that the "signal" provided to students may act a self-fulfilling prophecy: "*If the systems says I might drop out then maybe I should*". What are the safeguards that an institution can put around such 'intelligence' being misused?

Adaptive Learning Environments

While for many years researchers and educators have been interested in *Intelligent Tutoring Systems*, recently as the technology has matured and principles of adaptive learning have been incorporated into broader, more generic products and systems they have been labelled *adaptive learning* platforms. Adaptive learning platforms are among the most data-intensive learning tools, because they dynamically change the curriculum elements presented to each student, based on his or her history with the system. This is possible by comparing a model of a students' learning pathway and their understanding of concepts, with a model of what it means to master those concepts.

Adaptive learning environments are able to provide *fine-grained feedback* (e.g. which concepts have been grasped at what level), and *adaptive presentation* of sub-elements of the curriculum. Examples include presenting variations of a problem to check that a correct answer is indeed based on a robust understanding, backtracking after repeated failure on an assessment item to verify mastery of antecedent concepts, and not showing more advanced material that depends on mastery of concepts the learner has failed.

Examples of adaptive learning companies include Grockit, Knewton, Carnegie Learning, and Smart Sparrow, while publishers are also investing heavily in this area (e.g. Pearson and McGraw Hill Education). There are fewer examples of free tools to support adaptive learning compared with other forms of learning analytics driven tools, possibly due to the development complexity and cost.

Notable examples are the Open Learning Initiative based on research from Carnegie Mellon University, and ASSISTments. Khan Academy has also evolved from simple instructional videos into an adaptive platform.

While some may not consider them strictly adaptive learning environments, an emerging area is the use of natural language processing to support students with their writing. There is now mature research in the K-12 environment of coaching the basics of writing (e.g. Johnson et al., 2017) and commercial tutoring products are emerging (e.g. Turnitin's Revision Assistant; RedBird's Language Arts & Writing). At higher education levels, writing is obviously more advanced, and currently writing technology is at the stage of robust research platforms moving to large scale internal deployments (e.g. Buckingham Shum et al., 2016; Gibson et al., 2017; Knight, et al., 2018; Passonneau, et al., 2018).

At a more general level, new adaptive platforms have emerged in recent years to provide feedback to students based on data from a wide diversity of sources (see, for example, tools such as *OnTask*: Pardo, Jovanovic, Dawson, Gasevic & Mirriahi, 2019; Pardo et al., 2018; and *SRES*: Liu et al., 2017). Broadly speaking, these tools appropriate the techniques and technologies of "mass personalised messaging" (e.g. used by marketing companies), enabling an educator to send a "personal" email to hundreds of students, each one tuned, more or less, to the student's individual activity on a range of platforms. Each element of the message can be composed by the educator in their own voice, with encouragement and challenges based on their knowledge of the cohort. The email is then compiled dynamically by the tool, based on the degree to which the student has engaged in different activities, including attending lectures/labs, passing online quizzes, watching videos, posting to their learning blog, and so forth.

While there is evidence of the benefits of adaptive learning platforms and approaches (see Lovett et al., 2008), modelling the curriculum and the learner's understanding of it, and the process of structuring material for adaptive content engines is resource intensive. However, the expectation would be that any investment should pay off once hundreds or thousands of students are able to receive personalised feedback and adapted curriculum using the employed system.

The most advanced adaptive tools come from academic research programs in cognitive science and AI, which have modelled concept mastery and content adaptation techniques in great depth for a small number of domains. The material to be modelled must, by definition, have a structure and concepts that are robust, and assessment must be automated. These requirements, coupled with the academic grounding of the research teams in this field, sees most adaptive tools developed in the STEM disciplines. It is unclear to what extent arts, humanities and social sciences educators will benefit from adaptive learning platforms, with only a few examples reported to date (e.g. Ogan et al., 2006; Fournier-Viger et al., 2010).

Adaptive learning environments depend heavily on *models*: models of curriculum concepts, models of the learner's mastery of them (or of the learner themselves), models of how to teach concepts most effectively. Such models come with assumptions and values, which can be interrogated when they are made available to us, but they are always an incomplete lens on the true complexity of human learners and their environment.

A key ethical issue in this context is how a student's experience as a user within an adaptive learning environment should vary with their ability. What exactly will they see, or be blocked from seeing? As with predictive analytics, the accuracy – or integrity – of the models underpinning adaptive learning environments are critical. All the more so if these models run autonomously in the platform without the need for human intervention. The ability of an adaptive learning environment to capture or formatise curriculum, types of learners and learning, and assessment lies at the heart of how they are able to adequately "respond" to individual learners. Are there concerns with blocking students, based on performance, from certain curriculum materials? There are again ethical considerations about adaptive learning systems in misclassifying learners or implicitly labelling or diagnosing them, particularly if this "diagnosis" is then used to bar them from accessing educational materials available to other students.

Conversely there are ethical considerations associated with the role of teachers and teaching. In digital learning environments where tasks and feedback are adaptive and automated for students, what role does the teacher play? Are teachers "allowed" to assist when a student gets stuck? Who is responsible if the student is unable to pass a course after completing a core component of the curriculum within an adaptive learning environment? Moreover, are there ethical considerations associated with students' expectations to have a human teacher be responsible for instruction in core aspects of the curriculum?

Consistent with this, adaptive learning is certainly not without its critics. Some have argued that the rhetoric of "personalisation" that pervades the field of adaptive learning has overstated its educational value, detracting from the various strategies for personalisation that teachers already use, and from the value of peer discussion (see, for example, Feldstein and Hill (2014)). Healthy debate should be encouraged in this area, informed by a sound understanding of effective teaching practices, the limitations and capabilities of adaptive technologies, and clear perspective on data management practices. This should ward off partisan conversations either about the spectre of automation and AI or the panaceas offered by digital personalisation.

WiFi Usage

Equipment used to provide WiFi coverage contain the capacity to record information about their processes and the data used on the network. As part of this system, the address of the computers, mobile phones, tablets and any other devices that connect to the network can be uniquely identified. If the WiFi network being accessed requires authentication, the equipment also registers the credentials used for the access. With these two data items, the identity of a person connecting to the network can easily be obtained. This can then be combined with additional information such as the date and time of the event and the resources and/or services accessed. Whenever a device moves to a different location, the WiFi system can record its new location. The data collected by these systems can be easily processed to derive more comprehensive information.

The cost of collecting the information is virtually nothing. Tuning a set of devices to collect comprehensive data and provide it in a way that is suitable for further analysis requires additional effort, but usually doesn't represent an unreasonable cost. From the design and manufacturing point of view, there are clear benefits in capturing such a comprehensive account of what happens in individual digital devices; it can assist with an understand of circumstances that lead to a failure or anomalies, or it can be used to detect breaches and unauthorised access. The ramifications of this capacity in terms of privacy and security, however, are noteworthy.

Access to the information collected by these devices prompts numerous ethical concerns. Institutions need to carefully consider when it is necessary and appropriate to collect WiFi data and who will have access to this data. Location and access data could be useful for IT services to determine network load and requirements for infrastructure. Universities can use WiFi usage to determine, for example, the extent to which students are using the learning spaces of a University, and tailor their building and curriculum in light of this. Similarly, information on the websites, resources and services accessed could be used to ensure appropriate use of the network. However, in the learning context, it is vital that reasonable justification can be given as to why this kind of data is necessary to support learning processes. For example, instructors may design learning experiences that require students to be at a specific physical location during certain time (for example for a laboratory session), but this design decision should be decoupled from the need to gather evidence of such presence using information such as WiFi access. In circumstances where the collection of such data is deemed appropriate and necessary, it is important that students are made aware of the collection of such data and the purposes for which it will be used.

Cloud Storage of Data

Many universities are moving their IT architecture to the cloud. This means that computing services, such as software and databases, can be accessed from anywhere via an Internet connection. An increasing number of LMSs used in the Australian higher education context have moved to the cloud, including Canvas and Blackboard Ultra. Similarly, many universities are starting to make use of Customer Relationship Management services which run over the cloud (e.g. Salesforce) to provide students with a seamless service from pre-entry through to alumni status. Along with IT systems, much of the back-end supporting infrastructure is now provided over the cloud, with data warehouses and student and curriculum information systems now often hosted off site, and accessed via a connection to the cloud.

Moving to the cloud is a sensible transition for many different reasons: solutions can be flexibly scaled, disaster recovery is easier, the cost of hardware is effectively outsourced to the provider, document versioning is easier to control, and the loss of devices need not result in the loss of materials and data. Organisations do not need to worry about things like how much storage they have, or how much power it is consuming; they simply make use of the products, which are hosted somewhere else.

While the cloud provides a wide range of efficiencies when provisioning web-based services, it also brings a number of dangers due to the loss of control. Many universities are assured that they will be given access to “all the data” when they transition to cloud-based solutions, only to find out that vendors can have a very different conception of what this entails, and of what data universities may actually want. Similarly, an emphasis upon security alone can have negative effects when it comes to usability, data privacy, or portability. This is a particularly important aspect to consider within the context of the new European General Data Protection Regulation (GDPR), which Australian universities must carefully consider given they accept European citizens as students.

In developing suitable data governance procedures to administer the use of cloud storage for teaching and learning systems, consideration must be given to issues such as privacy and security. Processes must be put in place to monitor the access and accuracy of the data to ensure that it is being kept and used in ethical ways. Students should also be made aware of the data being collected and housed in such systems and the purposes for which this data will be used.

Institutional Case Studies

As research and development of learning analytics has increased, many educational institutions have begun to implement such tools and processes into their everyday practice. In this section we explore the approaches taken by different institutions in addressing ethical issues in the implementation of learning analytics.

The Australian Context[1]

Within the field of learning analytics, Australian researchers have played a significant role in the establishment of the field through their pioneering learning analytics research and tool development. This has meant that several Australian institutions have been using learning analytics tools within their practice for a number of years, yet many are still only now establishing processes and policies to support this use. In preparing this report we sought input from several Australian universities about their approach to ensuring the ethical use of learning analytics in their institutions. These approaches vary, but all seek to integrate consideration of the ethical use of data for learning analytics with existing policies and guidelines for the handling of institutional data more broadly.

Some institutions began with investigations and recommendations for implementation before learning analytics systems were put in place. At Flinders University in South Australia the *Flinders University Roadmap for Learning Analytics* was drafted to provide guidance for the institution on how to successfully implement learning analytics based on current research and practice in the field. Similarly, the University of Melbourne formed a Learning Analytics Working Group in 2012 to explore the key issues academic staff identified relating to the implementation of learning analytics, a key recommendation of which related to the ethical use of student data. This led to the drafting of a paper on key ethical issues relating to learning analytics which was presented to chancellery. At the University of Sydney consultations were held with a wide range of different stakeholders including students, learning analytics experts, privacy experts, academic development staff, student support staff, and IT services to develop the *Principles for the use of University-held student personal information for learning analytics at The University of Sydney.* Subsequently a Learning Analytics Advisory Board (LAAB) was established which includes representatives from across a number of stakeholder groups to oversee ongoing strategy and policy developments relating to learning analytics.

Other institutions have also established new positions and/or committees to investigate and implement policies and guidelines relating to the ethical use of learning analytics. At the Queensland University of Technology reviews have been conducted of current protocols and communications and a Chief Data Officer has been recruited whose role includes overseeing the enterprise-wide consistent and ethical use of learning analytics. At RMIT in Melbourne, a Learning Analytics Division has been established as part of the institution's strategy for the development and implementation of learning analytics. In addition, an Information Governance Board (IGB) has been formed to oversee the operational frameworks, guidelines and policy processes of data used in the institution, while the Information Stewards Group (ISG) is responsible for the operational oversight of the processes and outcomes related to data governance.

---

[1] In preparing this overview of the Australian context we'd like to thank the following people for sharing their experiences: Kathryn Bartimote, Kim Blackmore, David Fulcher, Andrew Gibson, Ann Luzeckyj, Pablo Munguia, Pip Pattison, and Josua Pienaar.

Several other institutions have developed ethical policies and guidelines in parallel to the development and/or implementation of learning analytics systems at the institutional level. The University of Wollongong (UOW) was among the first in the country to address ethical and privacy issues relating to learning analytics as they transitioned their learning analytics from individual projects to an institution-wide implementation. The University has two key documents relating to the ethical use of learning analytics: 1) the learning analytics data use policy, and 2) guidelines for actioning learning analytics insights. Matters relating to these documents are overseen by the UOW Learning Analytics Governance Group. At Central Queensland University, policy around the ethical use of learning analytics has been developed during the roll out of their new Learner Academic Prediction System (LAPS). The focus has been on guidelines for the use and management of student data, as issues relating to privacy and access are already covered in existing university policies. The LAPS project board are also developing a framework for interventions based on LAPS to support stakeholders in determining access to data and support services offered.

The consideration of ethical principles in the design of learning analytics systems and the involvement of stakeholders during the design process is increasingly becoming more prevalent in the Australian context. At the University of Canberra a participatory design approach was adopted in the development of the University's InterFace dashboard. This system provides learning analytics data to students, unit convenors and university executive staff and incorporates ethical design features such as the use of pseudonyms to anonymise student data for privacy protection. As the system continues to evolve the project team continually review the ethical considerations identified and make necessary design changes. At the University of Technology Sydney (UTS) an approach called the Ethical Design Critique (EDC) has been used to identify ethical considerations during the design phase of learning analytics system development. Developed by the UTS Connected Intelligence Centre, the approach involves a 2-3 hour workshop in which key stakeholders are asked to review the proposed design with a focus on the ethical dimensions. The design team then responds to this feedback which requires workshop participants' approval before the system can be submitted for senior executive signoff. UTS are also designing a Learning Record Store that will be compliant with regulations, such as the EU General Data Protection Regulation, and includes a Personal Data store that enables students to take their data with them when they graduate in a form that can be used across different technological platforms (e.g. xAPI).

International Case Studies

Internationally, educational institutions have also taken a diverse range of approaches to ensuring the ethical use of learning analytics. These are often tailored to respond to institutional cultures as well as national/regional regulatory frameworks. While an examination of a wide range of these international approaches is beyond the scope of this paper, we have included a profile of two international universities as examples of good practice:

University of British Columbia

Dr Simon Bates, Associate-Provost, Teaching and Learning

The goal of the Learning Analytics (LA) project at the University of British Columbia (UBC) is to better understand and to improve the learning experiences of students through the collection and analysis of relevant data, leading to data-informed decisions about enhancement of learning contexts, activities, courses and programs. Learners must be involved as active agents in this process, and as collaborators and co-interpreters, not simply as passive recipients. All aspects of this analytics activity must be pursued in a manner that is sensitive to the ethical and privacy concerns inherent in collection, analysis and retention of this data.

At UBC, most learning analytics activity falls under an umbrella of quality assurance and enhancement, and thus is not subject to formal Behavioural Research Ethics principles and practices. However, an appropriate and consistent approach is still required, and we therefore brought together a high-level academic committee to propose institutional principles, policy and practice with respect to learning data. The committee has wide representation from undergraduate and graduate students, the Office of the University Council, Provosts' offices, Heads, faculty members, the Centre for Teaching, Learning & Technology, the Chief Information Officer, the Chief Data Officer, the Registrar and the Director of Ethics for Research Services.

Over the last year, the Committee has developed draft principles to guide UBC's approach, taking into consideration work completed by (amongst others) JISC, the IMS Global Learning Consortium, the University of Edinburgh, the UK Open University, and SoLAR. This work is occurring at a time when faculty members and students are increasingly concerned about being monitored, how data is being used, and in the context of widespread misuse of data by social media platforms and affiliates.

These draft principles have then been tested against a number of pilot projects currently underway, derived from proposals submitted by faculty members, to ensure that they are comprehensive and useful, prior to being finalised. The principles include:

1. Respect for persons: Data and its analysis can never automatically provide the whole picture about a learner's likelihood of success or capability in their studies, and as such will never solely be used to inform actions of consequence at an individual level, as this must always involve human and personal intervention. We recognise that trends, norms or grouping of learners may introduce or reinforce bias in learner, faculty or institutional perceptions and behaviors, and will actively work to recognise and minimise these. We will practice 'data minimisation', collecting only what data is necessary.
2. Learners as autonomous agents: Learners are key stakeholders in LA, and will be involved in the LA project and all activities as collaborators and co-interpreters. They have the right to access the data collected related to their learning and, to act on it and, if necessary, to verify it.
3. Responsibility: Information that LA may provide should be used and acted upon if feasible to do so. As learners, as educators and as an institution we have a responsibility to use and extract meaning from learning data for the benefit of learners.
4. Equity: We will use learning analytics to help all learners achieve their learning goals, in order to succeed and excel, not merely those who may be deemed at risk of failure.
5. Stewardship and privacy: Data will be stewarded (collected, stored, granted access to, deleted) so as to comply with privacy and security best practices, policies and legislation, including adhering to principles of data minimization and individual choice / consent to the extent possible.
6. Accountability and transparency: Governance of LA activities will be ethically conducted, aligned with institutional policy, strategy and values and will include all stakeholders. It will include acknowledging the possibility of unforeseen consequences and mechanisms for redress. We will be transparent in communicating how data is collected, what is collected and how it is used. We will regularly report back to and engage with stakeholders.

7. Evolving and dynamic: As the use of learning data in new ways will have impacts on current assumptions and practices, we will commit to an on-going process of review and refinement of approaches, policies and practices as necessary, including regularly engaging all stakeholders, particularly students.

These principles guide our approach and general direction of travel, but at a process level more is needed. Uptake and interest are increasing across the campus and more questions are being asked that require access to and aggregation of data in order to be able to answer them. With the support of various stakeholder groups, we are currently working to develop an "agile, but equally robust" version of behavioural ethics review which considers who has access to data, what other data it is to be combined with, for what purpose, under whose authority and for what duration.

University of Edinburgh

Professor Dragan Gasevic, Professor of Learning Analytics & Yi-Shan Tsai

The University of Edinburgh had initially started exploring learning analytics in the context of its MOOC initiative that started in 2012. This early exploration lead to various pilot studies on learning analytics in the following years. These early initiatives encountered challenges such as low usability of data, significant requirements of effort and data skills, diversities across Schools in terms of the use of digital data, issues around data protection, and inconsistent expectations from stakeholders. As a result, a task group was established in late 2016 to drive the development of a learning analytics policy to ensure effective and responsible adoption of learning analytics. The task group was led by the Chair in Learning Analytics and Informatics and included representatives from senior management, Colleges, Information Services, Academic Services, Student Systems, the Edinburgh University's Students Association, and learning analytics researchers.

The policy development involved two phases: 1) Principles and purposes document that offers an succinct and easy-to-understand outline of the ways learning analytics are used at the University; and 2) Detailed policy document that defines procedures for issues such as data governance, transparency, informed consent, ethics, privacy, and access in line with the principles and purposes defined in Phase 1. Since then, a range of communication and engagement activities were undertaken to facilitate consultations with relevant stakeholders at a broad scale. These activities include discussions at Senate, the Senate Learning and Teaching Committee (LTC), the Joint Senate and Court Knowledge Strategy Committee (KSC), meetings with the University Schools and Colleges, surveys and focus groups with students and teaching staff, opportunities for stakeholders to provide written submissions on a set of Principles and Purposes, and open consultation for the University members to comment on the policy. In May/June 2017, LTC and KSC approved the Principles and Purposes of the University's approach to learning analytics. During 2017 to 2018, the Task Group continued to develop the Phase 2 document, which was recently submitted for approval by LTC and KSC. This whole process of inclusive policy development has raised significant awareness of learning analytics across the University and catalysed interest of a wide range of stakeholders to set out several new learning analytics initiatives at the University.

The LTC and KSC have established a review group to scrutinise plans for substantial new learning analytics activities. The group will assist proposers to align their activities with the University's Principles and to meet practical and regulatory issues (e.g. regarding data protection and security), and will also assist the University to share good practices in this developing area.

Conclusion

As outlined above, the goal of this paper has been to outline and expose key ethical issues associated with the use of learning analytics and data in digital learning environments. The paper has provided a brief introduction to learning analytics and general ethics considerations before proposing eight ethical principles that emerged from the literature in the area. It has reviewed key learning analytics ethical frameworks that already exist and has discussed in detail five common uses of learning analytics for which institutions are increasingly likely to need to have an ethical response. Finally, it has reviewed a number of institutional practices that are setting the benchmark in the ethical use of learning analytics.

The goal of this paper is not provide a set of firm recommendations for individuals, institutions or the higher education sector. We feel that this is not our place and, moreover, this is difficult to do in an emerging area of research and practice where tailored, localised responses are more likely to be appropriate. However, we do feel it would be useful to conclude the paper with a number of key considerations that have emerged for us in preparing this paper. These considerations are proposed tentatively below, but are also put forward as something of a "call to arms". They are intentionally presented and phrased in a way that might provoke educational leaders' and practitioners' to think about *what they might need to do* when it comes to the ethical use of learning analytics within their institutions.

We need to:

- **Recognise that the ethics of learning analytics is very complex**. While seemingly obvious, it is important to acknowledge the diverse ethical principles that are at play – we have proposed eight – and be mindful of the ways in which administrative, educational and learning analytics platforms and systems generate data that intersect with each of these principles. There is a need to recognise how data can and are already being used, and start interrogating these uses more closely with a structured ethical lens.

- **Develop clear principles and guidelines on data use in learning and teaching**. When one considers what is standard operating procedure for the use of data in research, it is clear how much work needs to be completed in many institutions before they might claim to have similarly robust principles and procedures for the ethical use of data in the area of learning and teaching. It is clear from our preparation of this paper that these principles and guidelines need to be established and need to sit outside the technology systems and applications that are used by institutions and individuals. In short, the technological systems employed should not set the principles of what data is collected and used within the institution.

- **Actively engage with multiple stakeholders**. There needs to be a "community-based" conversation about the ethical use of learning analytics and it is critically important to include multiple stakeholders, particularly students and staff, in this conversation. Such conversations will or should provide the dual purpose of sharing knowledge and information about the ethics of learning analytics while also promoting a shared communal understanding of the relevant issues and what should be developed and enacted. Such conversations could also consider how the university community can engage in advocacy about the design of educational technology and learning analytics systems that are are being developed internally and those developed commercially external to the university.

- **Establish transparency and trust.** Developing clear principles and genuinely engaging with stakeholders will go some way towards establishing transparency and trust across institutions in the area of learning analytics. The reality is that compared to more conventional technologies, advanced technologies that rely on AI or big data are often complex and difficult for non-technical people to understand. As we increasingly use advanced technologies that we do not understand all the time in our daily lives, system opacity might not appear to be a problem *per se*. However, we entrust our lives to aerospace engineers, surgeons and financial advisors knowing that they are regulated professionally, and we have experienced first-hand the benefits. In contrast, advanced educational technology systems that are difficult to understand are a much newer phenomenon: there has yet to be a generation of children or university students who have been immersed in such environments through to graduation. This highlights how important it is that stakeholders (e.g. educators, students, parents, instructional designers, policy makers) can *trust* the tool, and moreover, that educational institutions have suitably qualified staff asking probing questions of learning analytics and adaptive learning systems and providers. There is a need to establish or ensure transparency in whatever we do in this area. Advanced educational technologies that employ complex learning analytics and AI should not regarded as "Black Box" tools to be trusted but should be transparent and able to be challenged.

- **Avoid reinventing the wheel.** What is clear from the reviews undertaken in the preparation of this paper is that while discussion, research and practice in the area of the ethics of learning analytics is a relatively new, there are already well-established models, systems, principles and frameworks to draw on. We have highlighted a couple, but there are more. Institutions can, therefore, build on this foundation when considering their ethical responses to the use of learning analytics. It is also worth considering whether in addition to individual institutional responses to the issues raised in this paper, there is value in a sector wide response or manifesto.

- **Get a move on.** Higher education institutions, leaders and educational technology experts are in danger of being accused of burying their heads in the sand when it comes to the ethics of learning analytics. Experienced practitioners and researchers will be well aware that the issue of ethics in the field of learning analytics has been "burning" for some time. The Cambridge Analytica scandal (among others), the regular news reports of data breaches, and the GDPR are yet more reminders that, aside from everything else, the ethical use of students' and other data is a legal issue for universities. There is a need for action in order to prevent this area becoming a significant legal risk for institutions.

- **Develop processes to revisit and recast practice**. While there is clearly value in having well-crafted principles, policies and guidelines within an institution, it is also clear that as technology systems and approaches to data analytics evolve and change, new approaches to practice may need to be developed. It is essential that institutions do not adopt a "set and forget" mentality to the ethics of learning analytics as there will inevitably be a need to review and update policy and practice in this fast-moving area.

These seven considerations are not comprehensive, but we present them to trigger conversations about what action can be taken and what can be done in institutions and across the sector. Fundamentally, many of the questions and issues both institutions and individual face in the ethical use of learning analytics are about data governance, management and use.

While not intending to simplify the complexity of this challenge, this may be reduced to some key questions:

- Who has access?
- To what data?
- To do what?
- For what reason?
- And what has been learnt from this?

We hope this discussion paper provides a useful contribution to the ongoing conversation about a critical issue in higher education in Australia.

References

Arnold, K. E., & Sclater, N. (2017). Student perceptions of their privacy in learning analytics applications. In *Proceedings of the seventh international learning analytics & knowledge conference* (pp. 66-69). ACM.

Australian Human Rights Commission (2018). *Human Rights and Technology Issues Paper.* https://www.humanrights.gov.au/our-work/rights-and-freedoms/publications/human-rights-and-technology-issues-paper-2018

Bodily, R., Kay, J., Aleven, V., Jivet, I., Davis, D., Xhakaj, F., Verbert, K. (2018). Open Learner Models and Learning Analytics Dashboards: A Systematic Review. In *8th International Learning Analytics & Knowledge Conference (LAK 2018)*, Sydney: ACM Press.

Brooker, A., Corrin, L., Mirriahi, N. & Fisher, J. (2017). Defining "data" in conversations with students about the ethical use of learning analytics. In H. Partridge, K. Davis, & J. Thomas (Eds.), *Me, Us, IT! Proceedings ASCILITE2017: 34th International Conference on Innovation, Practice and Research in the Use of Educational Technologies in Tertiary Education* (pp. 27-31). ASCILITE.

Buckingham Shum, S. (2012). *Learning Analytics*. Policy Brief. UNESCO Institute for Information Technologies in Education. https://iite.unesco.org/publications/3214711/

Buckingham Shum, S. B., Knight, S., McNamara, D., Allen, L., Bektik, D., & Crossley, S. (2016). Critical perspectives on writing analytics. In *Proceedings of the sixth international conference on learning analytics & knowledge* (pp. 481-483). ACM.

Colvin, C., Rogers, T., Wade, A., Dawson, S., Gasevic, D., Buckingham Shum, S., Nelson, K., Alexander, S., Lockyer, L., Kennedy, G., Corrin, L., & Fisher, J. (2015). *Student retention and learning analytics: A snapshot of Australian practices and a framework for advancement*. Sydney: Office for Learning and Teaching.

Cormack, A. N. (2016). A data protection framework for learning analytics. *Journal of Learning Analytics, 3*(1), 91-106.

Corrin, L., & de Barba, P. (2014). Exploring students' interpretation of feedback delivered through learning analytics dashboards. In B. Hegarty, J. McDonald, & S.-K. Loke (Eds.), *Rhetoric and Reality: Critical perspectives on educational technology. Proceedings Ascilite, Dunedin 2014* (pp.629-633).

Corrin, L. (2019). Supporting the use of student-facing learning analytics in the classroom. In J. Lodge, J. Horvath & L. Corrin (Eds.), *From data and analytics to the classroom.* Routledge.

Corrin, L., Kennedy, G., de Barba, P. G., Lockyer, L., Gaševic´, D., Williams, D., . . . Bakharia, A. (2016). Completing the Loop: Returning Meaningful Learning Analytic Data to Teachers. Retrieved from Sydney: http://melbourne-cshe.unimelb.edu.au/__data/assets/pdf_file/0006/2083938/Loop_Handbook.pdf

Data & Society Institute (2018). *Algorithmic Accountability: A Primer.* https://datasociety.net/output/algorithmic-accountability-a-primer

Drachsler, H., Hoel, T., Cooper, A., Kismihók, G., Berg, A., Scheffel, M., Chen, W., & Ferguson, R. (2016). Ethical and privacy issues in the design of learning analytics applications. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge* (LAK '16). (pp. 492-493). ACM.

Drachsler, H., & Greller, W. (2016). Privacy and analytics: it's a DELICATE issue a checklist for trusted learning analytics. In *Proceedings of the sixth international conference on learning analytics & knowledge* (pp. 89-98). ACM.

Farah, J. C., Vozniuk, A., Rodriguez Triana, M. J., & Gillet, D. (2017). A Teacher Survey on Educational Data Management Practices: Tracking and Storage of Activity Traces. In *10th Workshop on Ethics & Privacy in Learning Analytics@ EC-TEL 2017* (No. EPFL-TALK-231023).

Feldstein, M., & Hill, P. (2016). Personalized learning: What it really is and why it really matters. *EDUCAUSE Review, 51*(2), 24-35.

Ferguson, R., Hoel, T., Scheffel, M., & Drachsler, H. (2016). Guest editorial: Ethics and privacy in learning analytics. *Journal of Learning Analytics, 3*(1), 5–15.

Fournier-Viger, P., Nkambou, R., & Nguifo, E. M. (2010). Building intelligent tutoring systems for ill-defined domains. In R. Nkambou, J. Bourdeau and R. Mizoguchi (Eds.) *Advances in intelligent tutoring systems* (pp. 81-101). Springer: Berlin, Heidelberg.

Gašević, D., Dawson, S., Rogers, T., & Gašević, D. (2016). Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting academic success. *The Internet and Higher Education, 28*, 68-84.

Gibson, A., Aitken, A., Sándor, Á., Buckingham Shum, S., Tsingos-Lucas, C., & Knight, S. (2017, March). Reflective writing analytics for actionable feedback. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference* (pp. 153-162). ACM.

Greller, W., & Drachsler, H. (2012). Translating learning into numbers: a generic framework for learning analytics. *Educational Technology & Society, 15*(3), 42-58.

Griffiths, D., Drachsler, H., Kickmeier-Rust, M., Hoel, T., & Greller, W. (2016). Is privacy a showstopper for learning analytics? A review of current issues and solutions. *Learning Analytics Review 6*. LACE. Retrieved from http://www.laceproject.eu/learning-analytics-review/is-privacy-a-show-stopper/

Gursoy, M. E., Inan, A., Nergiz, M. E., & Saygin, Y. (2017). Privacy-preserving learning analytics: Challenges and techniques. *IEEE Transactions on Learning technologies*, *10*(1), 68-81.

Heath, J. (2014). Contemporary privacy theory contributions to learning analytics. *Journal of Learning Analytics, 1*(1), 140-149.

Ifenthaler, D., & Tracey, M. W. (2016). Exploring the relationship of ethics and privacy in learning analytics and design: implications for the field of educational technology. *Educational Technology Research and Development*, *64*(5), 877-880.

Ifenthaler, D., & Schumacher, C. (2016). Student perceptions of privacy principles for learning analytics. *Educational Technology Research and Development*, *64*(5), 923-938.

Jivet, I., Scheffel, M., Specht, M., & Drachsler, H. (2018). License to evaluate: Preparing learning analytics dashboards for educational practice. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge* (pp. 31-40). ACM.

Johnson, A. M., McCarthy, K. S., Kopp, K. J., Perret, C. A., & McNamara, D. S. (2017). Adaptive Reading and Writing Instruction in iSTART and W-Pal. *FLAIRS 2017 - Proceedings of the 30th International Florida Artificial Intelligence Research Society Conference* (pp. 561-566). AAAI Press, Menlo Park.

Kitto, K., Williams, C., & Alderman, L. (2019). Beyond Average: Contemporary statistical techniques for analysing student evaluations of teaching. *Assessment & Evaluation in Higher Education, 44*(3), 338-360.

Knight, S., Shum, S. B., Ryan, P., Sándor, Á., & Wang, X. (2018). Designing academic writing analytics for civil law student self-assessment. *International Journal of Artificial Intelligence in Education, 28*(1), 1-28.

Kuzilek, J., Hlosta, M., Herrmannova, D., Zdrahal, Z., & Wolff, A. (2015). OU Analyse: analysing at-risk students at The Open University. *Learning Analytics Review*, LAK15 -1*,* 1-16.

Lang, C., Macfadyen, L. P., Slade, S., Prinsloo, P., & Sclater, N. (2018). The complexities of developing a personal code of ethics for learning analytics practitioners: Implications for institutions and the field. In *Proceedings of the 8th international conference on learning analytics and knowledge* (pp. 436-440). ACM.

Liu, D. Y. T., Bartimote-Aufflick, K., Pardo, A., & Bridgeman, A. J. (2017). Data-driven personalization of student learning support in higher education. In *Learning analytics: Fundaments, applications, and trends* (pp. 143-169). Springer, Cham.

Lovett, M., Meyer, O. & Thille, C., (2008). JIME - The Open Learning Initiative: Measuring the Effectiveness of the OLI Statistics Course in Accelerating Student Learning. *Journal of Interactive Media in Education*. 2008(1), p.Art. 13. https://www-jime.open.ac.uk/articles/10.5334/2008-14/galley/352/download/

Ogan, A., Aleven, V. and Jones, C. (2006). Culture in the classroom: Challenges for assessment in ill-defined domains. *Proc. Workshop on Intelligent Tutoring Systems for Ill-Defined Domains*, International Conference on Intelligent Tutoring Systems (pp. 92-100).

Pardo, A., & Siemens, G. (2014). Ethical and privacy principles for learning analytics. *British Journal of Educational Technology, 45*, 438-450. https://doi.org/10.1111/bjet.12152

Pardo, A., Bartimote-Aufflick, K., Buckingham Shum, S., Dawson, S., Gao, J., Gašević, D., ... Vigentini, L. (2018). OnTask: Delivering Data-Informed Personalized Learning Support Actions. *Journal of Learning Analytics, 5*(3), 235-249.

Pardo, A., Jovanović, J., Dawson, S., Gašević, D., & Mirriahi, N. (2019). Using Learning Analytics to Scale the Provision of Personalised Feedback. *British Journal of Educational Technology, 50*(1), 128-138.

Park, Y., & Jo, I. H. (2015). Development of the Learning Analytics Dashboard to Support

Students' Learning Performance. *Journal of Universal Computer Science, 21*(1), 110-133.

Passonneau, R. J., McNamara, D., Muresan, S., & Perin, D. (2017). Preface: special issue on multidisciplinary approaches to AI and education for reading and writing. International *Journal of Artificial Intelligence in Education, 27*(4), 665-670.

Prinsloo, P., & Slade, S. (2016). Student vulnerability, agency, and learning analytics: An exploration. *Journal of Learning Analytics*, *3*(1), 159-182.

Rienties, B., & Toetenel, L. (2016). The impact of learning design on student behaviour, satisfaction and performance: A cross-institutional comparison across 151 modules. *Computers in Human Behavior, 60*, 333-341.

Roberts, L. D., Howell, J. A., Seaman, K., & Gibson, D. C. (2016). Student Attitudes toward Learning Analytics in Higher Education:“the fitbit version of the learning world”. *Frontiers in Psychology*, 7, pp.1-11.

Roberts, L. D., Howell, J. A., & Seaman, K. (2017). Give me a customizable dashboard: Personalized learning analytics dashboards in higher education. *Technology, Knowledge and Learning*, 22(3), 317-333.

Rodríguez-Triana, M. J., Martínez-Monés, A., & Villagrá-Sobrino, S. (2016). Learning Analytics in Small-Scale Teacher-Led Innovations: Ethical and Data Privacy Issues. *Journal of Learning Analytics, 3*(1), 43-65.

Rubel, A., & Jones, K. M. (2016). Student privacy in learning analytics: An information ethics perspective. *The Information Society, 32*(2), 143-159.

Sclater, N. (2016). Developing a code of practice for learning analytics. *Journal of Learning Analytics, 3*(1), 16-42.

Sclater, N., & Bailey, P. (2015). JISC *Code of practice for learning analytics*. Available at: https://www.jisc.ac.uk/guides/code-of-practice-for-learning-analytics

Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist, 57*(10), 1380-1400.

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist, 57*(10), 1510-1529.

Steiner, C. M., Kickmeier-Rust, M. D., & Albert, D. (2016). LEA in private: a privacy and data protection framework for a learning analytics toolbox. *Journal of Learning Analytics*, *3*(1), 66-90.

Swenson, J., & Duin, A. H. (2018). Understanding Ethical Concerns Surrounding Big Data and Learning Analytics in Higher Education. *The Wiley Handbook of Educational Policy*, 479-511.

The Open University UK (2014). *Ethical use of Student Data for Learning Analytics.* Student Policies and Regulations. https://help.open.ac.uk/documents/policies/ethical-use-of-student-data

Teasley, S. D. (2017). Student facing dashboards: One size fits all?. *Technology, Knowledge and Learning, 22*(3), 377-384.

Teasley, S. D., & Whitmer, J. (2018). *The impact of student-facing LMS dashboards*. Available from: https://www.blackboard.com/resources/pdf/report-student-facingdashboards-finalhq.pdf

Turilli, M., & Floridi, L. (2009). The ethics of information transparency. *Ethics and Information Technology*, *11*(2), 105-112.

West, D., Huijser, H., Lizzio, A., Toohey, D., Miles, C., Searle, B., & Bronnimann, J. (2016a). *Learning Analytics: Assisting Universities with Student Retention, Final Report (Part 1)*. Sydney: Office for Learning and Teaching.

West, D., Huijser, H. & Heath, D. (2016b) Putting an ethical lens on learning analytics, *Education Technology Research and Development*, 64: 903-922.

Willis, J. E., Slade, S., & Prinsloo, P. (2016). Ethical oversight of student data in learning analytics: a typology derived from a cross-continental, cross-institutional perspective. *Educational Technology Research and Development, 64*(5), 881-901.