Chapter 28: Unpacking Student Privacy

Elana Zeide

Center for Information Technology Policy, Princeton University, USA
Information Society Project, Yale Law School, USA
Information Law Institute, New York University School of Law, USA

DOI: 10.18608/hla17.028

ABSTRACT

The learning analytics and education data mining discussed in this handbook hold great promise. At the same time, they raise important concerns about security, privacy, and the broader consequences of big data-driven education. This chapter describes the regulatory framework governing student data, its neglect of learning analytics and educational data mining, and proactive approaches to privacy. It is less about conveying specific rules and more about relevant concerns and solutions. Traditional student privacy law focuses on ensuring that parents or schools approve disclosure of student information. They are designed, however, to apply to paper "education records," not "student data." As a result, they no longer provide meaningful oversight. The primary federal student privacy statute does not even impose direct consequences for noncompliance or cover "learner" data collected directly from students. Newer privacy protections are uncoordinated, often prohibiting specific practices to disastrous effect or trying to limit "commercial" use. These also neglect the nuanced ethical issues that exist even when big data serves educational purposes. I propose a proactive approach that goes beyond mere compliance and includes explicitly considering broader consequences and ethics, putting explicit review protocols in place, providing meaningful transparency, and ensuring algorithmic accountability.

Keywords: MOOCs, virtual learning environments, personalized learning, student privacy, education data, ed tech, FERPA, SOPIPA, data ethics, research ethics

First, a caveat: the descriptions in this chapter should not be used as a guide to compliance, since legal requirements are constantly changing. It instead provides ways to think about the issues that people discuss under the banner of "student privacy" and the broader issues often neglected. In the United States, privacy rules vary across sectors. Traditional approaches to student privacy, most notably the Family Educational Rights and Privacy Act (FERPA),[1] rely on regulating how schools share and allow access to personally identifiable student information maintained in education records. They use informed consent and institutional oversight over data disclosure as a means to ensure that only actors with legitimate educational interests can access personally identifiable student information. This approach aligns with the Fair Information Practice Principles – typically notice, choice, access, and right to amend – that have been at the core of most privacy regulation since the early 1970s. These early privacy rules also focus primarily on disclosure of student information without addressing educators' collection, use, or retention of education records.

Newer approaches to student privacy tend to simply prohibit certain practices or require them to serve "educational" purposes. Blunt prohibitions are often crafted too crudely to work within the existing education data ecosystem, let alone support growth and innovation. "Education" purpose restrictions may limit explicit "commercial" use of student data, but they do not deal with the more nuanced issues raised by learning analytics and educational data mining even when used by educators for educational purposes. They do not consider the ways that using big data to serve education may not serve the interests of all educational stakeholders. It is difficult for categorically prohibitive legislation to be sufficiently flexible to match

[1] Family Educational Rights and Privacy Act (2014): see https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html?src=rn and https://www2.ed.gov/policy/gen/guid/fpco/pdf/ferparegs.pdf

the fast pace of technological change and the highly contextualized decision making in learning spaces. Data scientists and decision makers using learning analytics and education data mining must go beyond mere compliance through deliberate foresight, transparency, and accountability to ensure that data-driven tools achieve their goals, benefit the education system, and promote equity in broader society.

EDUCATION RECORD PRIVACY

The first wave of student privacy panic occurred in the late 1960s and early 1970s. Schools began to collect a wider array of information about students. Educators and administrators routinely shared student information on an ad hoc and often undocumented basis (Divoky, 1974).

FERPA's Default against School Disclosure

In response, Congress passed the primary federal statute governing student data, FERPA, in 1974. FERPA gives three rights to parents and "eligible students" over 18 or enrolled in postsecondary education ("parents," as shorthand). Federally funded schools, districts, and state education agencies must provide parents with access to education records maintained by the education institution or agency ("education actors," as shorthand) and the ability to challenge their accuracy. Education actors must also get parents' permission before sharing personally identifiable student information, subject to many exceptions that allow schools to consent on their behalf.[2]

FERPA focuses on limiting the disclosure of personally identifiable student information by educational institutions and agencies to approved recipients with legitimate educational interests. To meet FERPA's requirements, schools must obtain parents' written consent before sharing personally identifiable information maintained in a student's educational record unless one of several exceptions applies. In practice, the exceptions swallow the rule, and educators, not parents or students, make most privacy decisions (Zeide, 2016a).

Schools Authorizing Disclosure to Serve "Educational" Interests

The school official exception delegates the bulk of data-related decision making to schools and districts. Schools can share student personally identifiable student information without prior consent if the recipient is 1) performing services on their behalf and 2) has a "legitimate educational interest" in accessing such information; and, ostensibly, 3) has taken reasonable measures to exercise direct control over the information.[3] Educators decide what qualifies someone to be a school official and what constitutes a legitimate educational interest, but do not have to define these terms in any substantive detail (US Department of Education, n.d.). As a result, most rely on criteria so broad as to encompass almost any circumstance (Zeide, 2016a). Schools rarely take active measures to control recipients' detailed information practices, relying instead on terms of service or contracts between the parties as the means of "direct control" (Reidenberg et al., 2013).

Researchers Barred from Repurposing Student Data

FERPA places more stringent requirements on how educational actors share information with researchers. Under the studies exception, they must do so pursuant to a written contract with specific terms. Studies must be for the purpose of "developing, validating, or administering predictive tests; administering student aid programs; or improving instruction."[4] Researchers may only use personally identifiable student information for specified purposes and destroy the data once it is no longer needed.

Compliance-Oriented Enforcement

Educational actors have no direct accountability for FERPA violations. The statute is about putting a structure into place rather than preventing specific privacy violations. As a result, it does not impose consequences for individual instances of noncompliance. Students and educators cannot sue for violations under the statute (US Supreme Court, 2002). Instead, the US Department of Education (ED) has the power to withdraw all federal funding, including support in the form of federal student loans, if an educational institution or agency has a "policy or practice" of noncompliance.[5] However, the Department has never taken this dramatic action since the statute's enactment over forty years ago (Zeide, 2016b; Daggett, 2008). Since such a drastic measure would hurt the very students FERPA seeks to protect, the agency instead focuses on bringing education institutions into compliance. It is unlikely ED will ever pursue such a "nuclear" option (Solove, 2012).

STUDENT DATA PRIVACY

For almost forty years, stakeholders predominantly accepted FERPA's protection as sufficient despite minimal transparency, individual control over information, or consequences for specific violations in practice. FERPA's regulatory mechanisms no longer provide

[2] Family Education Rights and Privacy Act (FERPA), 20 U.S.C. § 1232g; 34 CFR Part 99 (2014); 34 CFR § 99.31 (exceptions): https://www.law.cornell.edu/uscode/text/20/1232g; https://www2.ed.gov/policy/gen/guid/fpco/pdf/ferparegs.pdf.

[3] Id. § 99.31(a)(1) (School Official Exception).

[4] Id. § 99.31(a)(6) (Studies Exception).

[5] 20 U.S.C. § 1232g(b)(1)–(2) (Policies or Practice Provision).

sufficient reassurance for stakeholders, in part because it regulates education records, not student data. The statute provides only narrow protection in terms of the information it covers, the actions it relates to, and the entities to which it applies in an age of big data.

From Education Records to Student Data

Low-cost storage, instantaneous transfer over connected networks, and cloud-based servers create an unprecedented volume, velocity, and variety of "big data" (Mayer-Schönberger & Cukier, 2014). Student information no longer means paper "education records" locked away in school filing cabinets, but rather interoperable, instantly transferable data stored on cloud servers. Interactive educational tools and platforms generate more information about students with more detail than has previously been possible. Data-mined information from out-of-classroom sources, like school ID geolocation and social media, goes far beyond traditional expectations regarding education records (Alamuddin, Brown, & Kurzweil, 2016). Even when mined student information is publicly available, many stakeholders find the notion of systematic collection and analysis of student data unsettling (Watters, 2015). The automatic capture of clickstream-level data about students, the permeability of cloud computing networks, and the infinite utility of big data prompts new privacy concerns (Singer, 2013).

FERPA's reliance on parental, student, or school oversight of recipients' information practices may not be possible, let alone practical or meaningful, given the quantity and complexity of big data and the automated transmission of information in interactive, digitally mediated environments. The statute does not even address schools' own privacy practices or cover new independent education providers, like massive open online courses (MOOCs), which collect information directly from users in "learning environments" but receive no federal funding. Stakeholders have little idea about what information schools and companies collect on students and how they use them (Barnes, 2014). They can't be sure that educators and data recipients even adhere to the privacy promises they make – especially when FERPA imposes no direct accountability for non-compliance.

Proliferation of New Student Privacy Protections

Since 2013, state policymakers responded to stakeholder panic by introducing over 410 student privacy bills: 36 states have passed 73 of these into law. On the federal level, legislators proposed amendments to FERPA and bills that would directly regulate the companies and organizations receiving student information. The vast majority of protective measures apply to federally funded P–12 public schools, but there is no consensus about what concerns matter and what "student privacy" means. This is clear from the incredible variety of ways districts, researchers, institutions, companies, states, and federal policymakers propose to protect student data (Center for Democracy and Technology, 2016; DQC, 2016; Vance, 2016).

Almost all reform measures reflect the need for more transparency, accountability, and baseline data safety, security, and governance protocols. Many simply continue FERPA's focus on how schools share information with third-party vendors and education researchers. Several explicitly prohibit school collection of certain types of information or from outside sources like social media. Some measures regulate data-reliant service providers directly (Center for Democracy and Technology, 2016; DQC, 2016; Vance, 2016).

Self-Regulation Supplements

More flexible approaches to privacy governance involve self-regulation. Over 300 companies have signed a Student Privacy Pledge,[6] created by the Future of Privacy Forum and the Software & Information Industry Association, which includes ten principles such as not selling student data. Signatories risk FTC enforcement if they do not abide by their promises (Singer, 2015). The US Department of Education, education organizations, and privacy experts are continuously releasing new best practice guidelines and privacy toolkits (Krueger, 2014; Privacy Technical Assistance Center, 2014). For stakeholders to have sufficient trust in these rules, however, there must be sufficient transparency about information practices, consideration regarding learning analytics purposes and potential outcomes, and accountability for noncompliance.

STUDENT PRIVACY GAPS

While the latest round of student privacy regulation has prompted much more explicit governance of student data and some sorely needed transparency, most reform measures still suffer from many of FERPA's flaws. Most students and stakeholders still have no concrete sense of what information is contained in education records, vague notions of how data can be used to their benefit, and minimal reassurance about what protections are in place (Prinsloo & Rowe, 2015; Rubel & Jones, 2016; Zeide, 2016a).

Minimal Meaningful Consent and Oversight

FERPA and similar rules rely on parental or school oversight of disclosure as a way to ensure that only appropriate recipients can access student data. This may not be possible, let alone practical or meaningful, given the quantity and complexity of big data and the automated transmission of information in interactive,

[6] http://studentprivacypledge.org/?page_id=45

digitally mediated environments. Contractual provisions create some more structure, but require schools to monitor third-party information practices and bring expensive lawsuits for enforcement. Finally, limiting disclosure doesn't work as well to prevent inappropriate use of student information since recipients who initially use this data for "legitimate education interests" often serve corporate or research interests at the same time (Young, 2015; Zeide, 2016a).

Crude Categorical Prohibitions

Some regulations attempt to address this problem by completely barring specific data collection, use, and repurposing. This often leads to problematic outcomes that conflict with current data use in the education system and unnecessarily restricts promising learning analytics and educational data mining. In Florida, for example, a ban on collecting biometric information conflicted with existing practices and legal obligations regarding special education students. As a result, many states have had to suspend or amend their initial attempts at ensuring student privacy. Erasure rules severely limit the potential for longitudinal studies and frequently often conflict with other record-keeping obligations imposed by state law (Vance, 2016).

Limits of Education Purpose Limitations

Many new laws follow the model of California's Student Online Personal Information Protection Act (SOPIPA), which covers entities providing education (K–12) services. They regulate how these actors use student information directly, rather than trying to do so through school oversight. Online providers of such services must have contracts with schools, erase student information upon request, and cannot create learner profiles that don't serve specified "K–12 purposes."[7]

Regulations that limit student data to "educational" use or purposes attempt to prevent commercial misuse by for-profit entities. Purpose limitations, however, do not address more nuanced issues raised by learning analytics and education data mining. Purpose limitation rules rest on the assumption of a consensus about what constitutes an "educational" purpose. They do not consider ways that institutions or researchers might prioritize goals other than the immediate educational interests of learner data subjects while still legitimately using data to manage institutional resources, improve the education system, or shed insight on learning science. Schools might, for example, use predictive data to exclude, rather than encourage, marginal students to save resources or improve rankings (Ashman et al., 2014; Drachsler & Greller, 2016; Rubel & Jones, 2016; Selwyn, 2014; Slade, 2016).

Leaving Out Learner Data

Most new laws do not address information held in higher education institutions. They do not address "learner" information collected by virtual learning environments independent of traditional, federally funded education institutions. Instead, the more permissive commercial privacy regime governs data collected and used by these private entities (in the absence of applicable state law). This means that use and disclosure of this "learner" data is limited by consumer privacy policies, which are notorious for being incomprehensible, overly broad, and open to change without notice (Jones & Regner, 2015; Polonetsky & Tene, 2014; Young, 2015; Zeide, 2016a).

EDUCATION DATA ETHICS

Society grapples with these issues across sectors, but they are particularly acute in education environments. As individuals who seek to improve education experiences and operate with integrity, it is easy to lose sight of how revolutionary the information practices involved in learning analytics and education data mining are compared to traditional education information practices and norms about student data. Students rarely have a realistic choice to opt out of mainstream data-driven technologies. Education data subjects are more vulnerable than those in typical consumer contexts, not only because they might be children, but also because learning requires some degree of risk-taking for intellectual growth. There are still unresolved issues about whether these tools may inadvertently reduce, rather than expand equitable opportunities, undermine the broader goals of the education system, and give students less agency and make them more, not less, vulnerable (Prinsloo & Slade, 2016; Siemens, 2014).

Equitable Outcomes

It is important for those working with student data to consider how consequences may play out in an inevitably flawed reality rather that the neutral space of theoretical and technological models. Algorithmic models may inadvertently discriminate against minorities or students of lower socioeconomic status. They may have disparate impacts. Tools that predict student success could repeat past inequities instead of promoting more achievement and upward mobility. Ostensibly neutral policies can create deeply inequitable outcomes due to uneven implementation (boyd & Crawford, 2011; Citron & Pasquale, 2014; Barocas & Selbst, 2014).

[7] The statute does, however, include a carve out indicating that it "does not limit the ability of an operator to use information, including covered information, for adaptive or personalized student learning purposes." § 22584(l). As of the writing of this chapter, it is not clear how these rules will work in practice. Student Online Personal Information Protection Act (SOPIPA), CAL. BUS. & PROF. CODE §§ 22584–22585 (2014), https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201320140SB1177.

Broader Education Effects

Continuously collecting detailed information in classrooms, from cameras, or from sensors can have broader consequences. Ubiquitous surveillance and embedded assessment may have a chilling effect on student participation and expression (Boninger & Molnar, 2016; Vance & Tucker, 2016). While these practices reduce reliance on periodic high-stakes tests, they also put every moment of the learning process under scrutiny. This may ultimately undermine trust in data-driven education tools and practitioners, chilling the intellectual risk-taking required in learning environments.

Inadvertently Shifting Authority

Learning analytics and educational data mining changes not only how, but who makes pedagogical and academic decisions. Traditionally, the individuals who evaluated and made decisions about students were close at hand and relied on personal, contextualized observation and knowledge. Parents, students, or administrators with concerns about particular outcomes could go directly to the relevant decision maker for explanation. This created transparency, and an easy avenue to seek redress, thereby providing accountability.

In adopting data-driven education tools, educators change what goes into measuring learning, what goals we seek to achieve through education, and who gets to make those decisions. Automated and algorithmic pedagogical and institutional decision-making shifts the locus of authority from a traditional, physically present human to obscure technologies or remote companies and researchers. Data-driven education changes who gets to make important decisions that shape lives and the education system overall. It does so without the shift being obvious, and, in many cases, deliberate. This shift in who can access and use data shifts power relationships as well. As security expert Bruce Schneier (2008) notes, "Who controls our data controls our lives" (paragraph 5). We must explicitly consider the handoff of authority that goes with the handoff of data.

GOING BEYOND COMPLIANCE

Under the current and emerging regulatory framework, learning analytics and education data mining practitioners and consumers will have much of that power. They will accordingly bear the responsibility of defining what student privacy means. Their decisions about technological structures, conceptual models, and learning outcomes craft the rules that apply in practice to information in learning environments. These decisions need to be made thoughtfully and deliberately. It also benefits learning analytics and educational data mining as a field by cultivating the trust required for individual participation, institutional implementation, and policymaker support for learning analytics and educational data mining overall.

I recommend going beyond mere compliance to take a more proactive approach. Ideally, this involves not only anticipating potential problems, but also putting protocols in place to determine practices if they arise and open communication with data subjects and stakeholders. Key components of proactive student privacy practices include 1) considering ethical implications; 2) creating explicit protocols for review; 3) actively communicating with data subjects and stakeholders about data practices, purposes, and protection; and 4) ensuring algorithmic accountability.

Ethical Scrutiny

Learning analytics and educational data mining projects should include deliberate, proactive consideration of potential benefits and their distribution across society and time, unintended outcomes related to learning and broader society, and ethical questions regarding experiment protocols and ultimate priorities. These reflect important considerations regarding human subject experiments promulgated in the Belmont Report in 1978 and later codified and institutionalized through Institutional Review Boards (IRBs) that must approve of academic research. However, data use only inside institutions, activities categorized as "optimization" instead of research, and company practices rarely undergo similarly explicit consideration of fundamental ethical principles.

Learning analytics and educational data mining practitioners, consortia, and supporters have promulgated ethical principles to guide information practices. These raise important issues, including the importance and difficulty of user notice and consent to how data is collected, stored, processed, and shared in learning systems, given the volume of information and complexity of algorithmic analysis. They also include more abstract notions of justice and beneficence that take into account whether experimental results serve the "greater good" (Drachsler & Greller, 2016; Open University, 2017; Pardo & Siemens, 2014; Sclater & Bailey, 2015; Slade, 2016; Asilomar, 2014).

Explicit Review

Privacy and ethical considerations should be incorporated from the first stages of technology and experimental design. At a minimum, data-driven education tools should be audited for unintended bias, disparate impact, and disproportionate distribution of risk and benefits across society. A best practice would create proactive measures to address possible, but foreseeable, problematic outcomes ahead of time. Is there a point, for example, when the discrepancy between two experimental groups is so high that researchers and educators should stop A/B testing?

Projects should have predetermined points for explicit accountability and ethical review. Many companies, for example, have begun to employ their own "consumer review boards" to take data subjects' and broader society's interests into account before moving forward on experiments and again before publication (Calo, 2013; Jackman & Kanerva, 2016; Tene & Polonetsky, 2015).

Aggressive Transparency

Ideally, learning analytics and education data mining tools and technologies should also provide meaningful transparency and algorithmic accountability. Transparency is important on both the micro- and the macro-level. Disclosing information practices helps reassure stakeholders who might panic in an absence of sufficiently specific and readily available information about learning analytics and educational data mining data practices.

Transparency and outreach about the ways that data analysis may benefit current learners – and not some future student in a land far, far, away – helps ameliorate stakeholder fears. Open and early communication also helps reduce the impression that a small elite group of scientists have tremendous control over student experiences and outcomes, and that their actions are shrouded in secrecy. It helps to recruit institutional resources to find ways to reach out to data subjects and the wider community.

Algorithmic Accountability

Transparency, however, is not enough to ensure appropriate information practices. It is a prerequisite. Documentation and accountability are also important given the stakes at issue and the obscurity of algorithmic decision making. Learners and stakeholders will want to know what evidence backs up pedagogical and institutional decision making. Ideally, learning analytics and education data mining practitioners should implement tools for algorithmic accountability. These include audits to double check that algorithmic tools perform as intended and actually promote promised outcomes.

A key piece of algorithmic accountability that will become increasingly important in affecting learners' future opportunities is the need to document algorithmic and institutional decision making to allow for due process (Diakopoulos, 2016; Kobie, 2016; Kroll et al., 2017). Learners, educators, and institutions will want to see the evidence and know about the systems that impact their academic progress and credentialing and examine the decisions that affect them (Zeide, 2016a; see also Citron & Pasquale, 2014; Crawford & Schultz, 2014). Data scientists and data-driven decision makers should be prepared to facilitate forensic examination of important decisions. Parents, for example, will want an explanation as to why their child was or was not promoted to the next grade. "Because the algorithm said so" will not be a sufficient response.

CONCLUSION

Trust is crucial to learning environments, which seek to foster intellectual experimentation and growth. As noted in a 2014 White House report on big data, "As learning itself is a process of trial and error, it is particularly important to use data in a manner that allows the benefits of those innovations, but still allows a safe space for students to explore, make mistakes, and learn without concern that there will be long term consequences for errors that are part of the learning process."

By going beyond mere compliance, those entrusted with education data can guard against potential unintended consequences of even the most well-meaning projects that might undermine the very goals they seek to achieve. The readers of this handbook entrusted with the wealth of student data should take a proactive approach that aims not at mere compliance, but goes beyond to consider broader social, ethical, and political implications. Doing so will promote trust in data-driven education and ensure that learning analytics and educational data mining achieve their revolutionary potential.

REFERENCES

Alamuddin, R., Brown, J., & Kurzweil, M. (2016). *Student data in the digital era: An overview of current practices*. Ithaka S+R. doi:10.18665/sr.283890

Ashman, H., Brailsford, T., Cristea, A. I., Sheng, Q. Z., Stewart, C., Toms, E. G., & Wade, V. (2014). The ethical and social implications of personalization technologies for e-learning. *Information & Management*, 51(6), 819–832. doi:10.1016/j.im.2014.04.003

Asilomar Convention for Learning Research in Higher Education. (2014). *Student data policy and data use messaging for consideration at Asilomar II*. Asilomar, CA. http://asilomar-highered.info/

Barnes, K. (2014, March 6). Why a "Student Privacy Bill of Rights" is desperately needed. Washington Post. http://www.washingtonpost.com/blogs/answer-sheet/wp/2014/03/06/why-a-student-privacy-bill-of-rights-is-desperately-needed/

Barocas, S., & Selbst, A. D. (2014). *Big data's disparate impact*. SSRN Scholarly Paper. Elsevier. http://papers.ssrn.com/abstract=2477899

Boninger, F., & Molnar, A. (2016). *Learning to be watched: Surveillance culture at school*. Boulder, CO: National Education Policy Center. http://nepc.colorado.edu/files/publications/RB%20Boninger-Molnar%20Trends.pdf

boyd, d., & Crawford, K. (2011). Six provocations for big data. SSRN Scholarly Paper. Elsevier. http://papers.ssrn.com/abstract=1926431

Calo, R. (2013). Consumer subject review boards: A thought experiment. *Stanford Law Review Online*, 66, 97.

Center for Democracy and Technology (2016, October 5). *State student privacy law compendium*. https://cdt.org/insight/state-student-privacy-law-compendium/

Citron, D. K., & Pasquale, F. A. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, 89. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2376209

Crawford, K., & Schultz, J. (2014). Big data and due process: Toward a framework to redress predictive privacy harms. *Boston College Law Review*, 55(1), 93.

Daggett, L. M. (2008). FERPA in the twenty-first century: Failure to effectively regulate privacy for all students. *Catholic University Law Review*, 58, 59.

Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56–62. doi:10.1145/2844110

Divoky, D. (1974, March 31). How secret school records can hurt your child. *Parade Magazine*, 4–5.

DQC (Data Quality Campaign) (2016). *Student data privacy legislation: A summary of 2016 state legislation*. http://2pido73em67o3eytaq1cp8au.wpengine.netdna-cdn.com/wp-content/uploads/2016/09/DQC-Legislative-summary-09232016.pdf

Drachsler, H., & Greller, W. (2016). Privacy and analytics – it's a DELICATE issue. A checklist to establish trusted learning analytics. Open Universiteit. http://dspace.ou.nl/bitstream/1820/6381/1/Privacy%20a%20DELICATE%20issue%20(Drachsler%20%26%20Greller)%20-%20submitted.pdf

Jackman, M., & Kanerva, L. (2016). Evolving the IRB: Building robust review for industry research. *Washington and Lee Law Review Online*, 72(3), 442.

Jones, M. L., & Regner, L. (2015, August 19). Users or students? Privacy in university MOOCS. *Science and Engineering Ethics*, 22(5), 1473–1496. doi:10.1007/s11948-015-9692-7

Kobie, N. (2016, January 29). Why algorithms need to be accountable. *Wired UK*. http://www.wired.co.uk/news/archive/2016-01/29/make-algorithms-accountable

Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165. http://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2765268

Krueger, K. R. (2014). 10 steps that protect the privacy of student data. *The Journal*, 41(6), 8.

Mayer-Schönberger, V., & Cukier, K. (2014). *Learning with big data*. Eamon Dolan/Houghton Mifflin Harcourt.

Open University. (2017). Ethical use of student data for learning analytics policy. http://www.open.ac.uk/students/charter/essential-documents/ethical-use-student-data-learning-analytics-policy

Pardo, A., & Siemens, G. (2014). Ethical and privacy principles for learning analytics. *British Journal of Educational Technology*, 45(3), 438–450. doi:10.1111/bjet.12152

Polonetsky, J., & Tene, O. (2014). Who is reading whom now: Privacy in education from books to MOOCs. *Vanderbilt Journal of Entertainment & Technology Law*, 17, 927.

Prinsloo, P., & Rowe, M. (2015). Ethical considerations in using student data in an era of "big data." In W. R. Kilfoil (Ed.), *Moving beyond the Hype: A Contextualised View of Learning with Technology in Higher Education* (pp. 59–64). Pretoria, South Africa: Universities South Africa.

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

Privacy Technical Assistance Center. (2014) *Protecting student privacy while using online educational services: Requirements and best practice*. http://ptac.ed.gov/sites/default/files/Student%20Privacy%20and%20Online%20Educational%20Services%20%28February%202014%29.pdf

Reidenberg, J. R., Russell, N. C., Kovnot, J., Norton, T. B., Cloutier, R., & Alvarado, D. (2013). *Privacy and cloud computing in public schools*. Bronx, NY: Fordham Center on Law and Information Policy. http://ir.lawnet.fordham.edu/clip/2/?utm_source=ir.lawnet.fordham.edu%2Fclip%2F2&utm_medium=PDF&utm_campaign=PDFCoverPages

Rubel, A., & Jones, K. M. L. (2016). Student privacy in learning analytics: An information ethics perspective. *The Information Society*, 32(2), 143–159. doi:10.1080/01972243.2016.1130502

Schneier, B. (2008, May 15). Our data, ourselves. *Wired Magazine*. http://www.wired.com/2008/05/security-matters-0515/all/1/

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

Selwyn, N. (2014). *Distrusting educational technology: Critical questions for changing times*. New York/London: Routledge, Taylor & Francis Group.

Siemens, G. (2014, January 13). The vulnerability of learning. http://www.elearnspace.org/blog/2014/01/13/the-vulnerability-of-learning/

Singer, N. (2015, March 5). Digital learning companies falling short of student privacy pledge. *New York Times*. http://bits.blogs.nytimes.com/2015/03/05/digital-learning-companies-falling-short-of-student-privacy-pledge/

Singer, N. (2013, December 13). Schools use web tools, and data is seen at risk. *New York Times*. http://www.nytimes.com/2013/12/13/education/schools-use-web-tools-and-data-is-seen-at-risk.html

Slade, S. (2016). Applications of student data in higher education: Issues and ethical considerations. Presented at Asilomar II: Student Data and Records in the Digital Era. Asilomar, CA.

Solove, D. J. (2012). FERPA and the cloud: Why FERPA desperately needs reform. http://www.law.nyu.edu/sites/default/files/ECM_PRO_074960.pdf

Tene, O., & Polonetsky, J. (2015). Beyond IRBs: Ethical guidelines for data research. *Beyond IRBs: Ethical review processes for big data research*. https://bigdata.fpf.org/wp-content/uploads/2015/12/Tene-Polonetsky-Beyond-IRBs-Ethical-Guidelines-for-Data-Research1.pdf

US Department of Education (n.d.). FERPA frequently asked questions: FERPA for school officials. Family Policy Compliance Office. http://familypolicy.ed.gov/faq-page/ferpa-school-officials

US Supreme Court. (2002). Gonzaga Univ. v. Doe, 536 U.S. 273. https://supreme.justia.com/cases/federal/us/536/273/case.html

Vance, A. (2016). *Policymaking on education data privacy: Lessons learned*. Education Leaders Report, 2(1). Alexandria, VA: National Association of State Boards of Education.

Vance, A., & Tucker, J. W. (2016). *School surveillance: The consequences for equity and privacy*. Education Leaders Report, 2(4). Alexandria, VA: National Association of State Boards of Education.

Watters, A. (2015, March 17). Pearson, PARCC, privacy, surveillance, & trust. Hack Education: The History of the Future of Education Technology. http://hackeducation.com/2015/03/17/pearson-spy

White House. (2014). *Big data: Seizing opportunities, preserving values*. Washington, DC: Exceutive Office of the President.

Young, E. (2015). Educational privacy in the online classroom: FERPA, MOOCs, and the big data conundrum. *Harvard Journal of Law and Technology*, *28*(2), 549–593.

Zeide, E. (2016a). Student privacy principles for the age of big data: Moving beyond FERPA and FIPPs. *Drexel Law Review*, 8(2), 339.

Zeide, E. (2016b, March 16). Interview with ED Chief Privacy Officer Kathleen Styles. Washington, D.C.