Received: 28 April 2022 | Accepted: 18 March 2023

DOI: 10.1111/bjet.13323

ORIGINAL ARTICLE

British Journal of Educational Technology BERA

A reflective design case of practical micro-ethics in learning analytics

Simon Knight[1,2] | **Antonette Shibani**[1] | **Simon Buckingham Shum**[2,3]

[1]TD School, University of Technology Sydney, Sydney, New South Wales, Australia

[2]Centre for Research on Education in a Digital Society (CREDS), Sydney, New South Wales, Australia

[3]Connected Intelligence Centre, University of Technology Sydney, Sydney, New South Wales, Australia

Correspondence

Simon Knight, TD School, University of Technology Sydney, and Centre for Research on Education in a Digital Society (CREDS), PO Box 123, Broadway, Ultimo, Sydney, NSW 2007, Australia.
Email: simon.knight@uts.edu.au

Abstract

A burgeoning literature, policy landscape, and set of practice resources and guidelines have emerged around ethical educational technology development and implementation, particularly in the context of data or artificial intelligence informed tools. However, while these resources provide valuable tools to support navigation of the ethical landscape, they suffer from some limitations. These include a focus on coarse-grained, clear cut cases; a lack of attention to dilemmas and tensions; and a potential focus on procedural aspects of ethics, rather than our moment-to-moment ethics-in-action. This paper provides a case study description, using a reflective design case to provide a more textured micro-ethics. The case approach, its exemplification as a tool for micro-ethics, and the features of interest in our particular case each provide valuable tools for the educational technology community.

KEYWORDS

case study, ethics, ethics-in-action, learning analytics, micro-ethics

INTRODUCTION

Recognition of ethical challenges in educational technology has existed for as long as edtech's potential has been discussed (Yeaman et al., 2013). To address these challenges in research, and specifically learning technology research, a rich set of material tools has been

Practitioner notes

What is already known about this topic

- Ethical concerns are central to the design, development and implementation of educational technologies (edtech).
- A range of guidelines, policies and law exist, providing principles that guide edtech research and use.
- However existing resources tend to reflect coarse procedural aspects of ethics, providing a navigation aid, but not a defined path, in ethical edtech practice.

What this paper adds

- This paper provides a novel method—a reflective design case—for describing and exploring micro-ethics in edtech.
- Through applying this method, we demonstrate that micro-ethical analyses of 'ethics-in-action' help to uncover the nuanced decisions we make in ethical conduct.
- This exemplification provides insight into a specific body of work on writing analytics, of relevance to the field.

Implications for practice and/or policy

- Policy and guidance for ethical conduct should make use of cases that demonstrate how such ethical materials are drawn on in ethical conduct, and as demonstrations of micro-ethics in action.
- Practitioners—those who design, develop, and implement edtech—may use the method developed to explore and share their own cases to support their work.
- Practitioners—those who design, develop, and implement edtech—may draw on shared reflective design cases in learning regarding navigating application of procedural ethics and developing ethics-in-action.

produced to support policy, research and practice, through a burgeoning range of principles, frameworks, and studies of user preferences, that position technologies such as artificial intelligence in education with respect to their ethical dimension and social context (Holmes et al., 2021; Kitto & Knight, 2019). Recent systematic review research has highlighted a range of practical supports and design principles for processes of design, development and implementation of AI tools across contexts (not with a focus on education) (Lu et al., 2022).

However, in everyday ethical action many issues of ethics do not fall into clear cut cases (Kitto & Knight, 2019). Ethical dilemmas may occur whereby incommensurable principles must be applied (such as privacy vs free speech), such that either action taken will result in some negative impacts. Perhaps more significantly, many everyday—even trivial—decisions have moral flavour, and we navigate these in the course of our interactions with each other and research processes (BERA, 2018). This ethical dimension to our decisions can be foreground by concrete examples (BERA, 2018). These help us discuss, and build capacity for addressing, ethical dilemmas, for example in considering the design, development and implementation of learning analytics systems (Kitto & Knight, 2019), or in considering how educators navigate the inherent dilemmas in assessment structures (Johnson et al., 2017).

There have therefore been calls for a more *textured* approach to ethics (Gray & Boling, 2016; Holmes et al., 2021; Kitto & Knight, 2019). That is, an ethics that both recognises the micro-interactional design contexts in which teachers, students and technologists make decisions that are imbued with ethical character, as well as the macro-situational context in which technologies act and interact, producing both hard (ie, direct) and soft (ie, indirect) impacts (Swierstra, 2015). In research ethics, these calls indicate a greater

articulation of the ethical concepts used (Beauchemin et al., 2021) and reflect a shift from formal or procedural ethics, towards a situational or *ethics-in-practice* approach (Guillemin & Gillam, 2004). This latter approach recognises both that the *application* of procedural ethics is a practice in its own right (applied ethics), but moreover a recognition that ethics are imbued in our action and design in ways not captured by procedural approaches, and the significance of sharing ethical cases and principles beyond *failures* of ethics. This shift is paralleled in moves towards 'micro-ethics', with recent work to integrate such approaches into data science education to connect "the 'big picture' of digital/data ethics to the routines of daily practice" (Bezuidenhout & Ratti, 2021, p. 939).

One lens onto the conceptual resources drawn on in ethical action is the material resources developed in the field of research and practice (for example, frameworks and policies). In addition, a limited range of concrete examples of ethical breaches or concerns have been presented both in research and coverage of concerns in news media. However, despite this popular awareness and set of resources, there are few practical examples or cases that provide close ethical readings of choices made in the design, development and delivery of educational technologies. For example, of the 3243 citations to Guillemin and Gillam's (2004) seminal paper on 'ethically important moments' in qualitative research, few of these papers have edtech as the central focus.[1] As such, relatively little is known about the micro-ethical (Guillemin & Gillam, 2004) decision making in the design, development and implementation of learning technologies.

That is therefore the aim of this paper; to provide an example *reflective design case* of ethics-in-action, and the granular micro-ethics inscribed and enacted in the design, development and implementation of a learning technology (a writing analytics tool). This distinctive approach provides a practical and nuanced reflexive response to situate ethics-in-action for educational technologies. The following sections thus correspond to the key contributions of this work:

1. In Section 2, *Method*, we outline a case-based approach to foreground design decisions at a micro level; addressing the methodological question "How can we support the shared description and exploration of micro-ethics in edtech?"
2. In Section 3, *Our Case* we implement this approach to demonstrate the potential of a micro-ethics approach, as expressed through our case method, addressing the theoretical question "What can a micro-ethics approach to edtech reveal regarding ethical practice?"
3. This exemplification foregrounds particular features of interest in our case for the field, addressing the case question *"What micro-ethical concerns arose in the design, development, and implementation of a writing analytics tool?"*

METHOD: THE REFLECTIVE DESIGN CASE

A reflective design case is used to draw attention to ethical features of the design, development and implementation of the learning technology under discussion. We explain its rationale in this section.

Our *case study* approach uses the *specific* case to demonstrate the ways in which a close analysis provides insight into the ethical texture of the learning technology. That is, the details of this specific case are not the principle focus (an intrinsic case study), but rather are intended to provide insights into an issue (microethics), through a systematic discussion of the issue in a particular context to provide an instrumental case (Stake, 2003). Insofar as the aim is to produce generalisations, it is to produce fuzzy generalisations of *method* (an approach to close ethical analysis), not of *outcome* (codified ethical decisions to be made).

We specifically adopt a *reflective case study* approach drawing on Guillemin and Gillam (2004) noting the importance of reflexivity in achieving ethical practice at a micro level. Reflexivity,[2] for them, means a sensitivity or attunement to "ethically important moments", and a responsiveness to these "difficult, often subtle, and usually unpredictable situations that arise in the practice of doing research" (Guillemin & Gillam, 2004, p. 262). These include features such as how to navigate participant disclosures, judgements made in pursuing lines of questioning that may be uncomfortable, and so on. Significantly, guidelines and principles cannot 'solve' the issue: "Research ethics committees cannot help when you are in the field and difficult, unexpected situations arise, when you are forced to make immediate decisions about ethical concerns" (Guillemin & Gillam, 2004, p. 273). Cases, though, provide a valuable tool for sharing and thinking about developing and applying ethical practice (Johansen & Frederiksen, 2021).

Our approach uses a *design case* to produce a *reflective design case* explicitly because design *practice* involves both design development *within* the bounds of theory, and the challenging of that theory (Sengers et al., 2005). Through 'reflection-in-action' (Schön, 1992) and our expressions of this in both design artefacts and deliberate reflective outputs, we probe and engage our application of theory in situational—and unanticipated—action (Sengers et al., 2005). Indeed, in a re-*phrasing* that provides a re-*framing* of Schön (1992) Frauenberger et al. (2017) suggest that: "in performing design and research, 'we show ourselves to be ethical [*orig*: *knowledgeable*] in a special way' (p. 49)—our ethics is tacitly *in action*, deeply folded into 'recognitions, judgements and skilful performances' (Schön, 1992, p. 50)." (Frauenberger et al., 2017, p. 11). As such, they propose that reflection-in-action is central to researchers' navigating the kinds of unexpected situations in the context of human-computer-interaction (HCI) research. They argue for the centrality of flexibility, and situational awareness or sensitivity to what we called "ethically important moments" above, beyond anticipatory planning (Bardzell et al., 2014).

Taking a *reflective design case* approach in the context of *learning* holds some distinct characteristics, compared to other HCI contexts. It suggests the necessity of understanding the *learning context* into which the designed artefacts are introduced, as mutually constitutive features of the learning task. As such, the reflective design case in the context of learning analytics builds on established work in design for learning, including Gray and Boling's (2016) analysis of 53 published design cases in the Journal of Designs for Learning, which documents rich design artefacts 'as they are' to describe design 'as it has taken place' and give insight into the implied or explicit values in these. Echoing the call for ethical cases above, these authors also call for more studies involving close analysis of design cases. Building on this background we adapt Bardzell et al.'s process (2014, p. 1955) to implement a design case of a learning technology. We concur with their assessment that design involves making judgements, and thus analysis of designs helps us engage in critical judgement regarding the purposes of design artefacts, whether it does or does not work (and for whom), and so on (Bardzell et al., 2014). Our approach differs insofar as their analysis is specifically intended to draw out *critical* (qua, critique) aspects of design or designs as critique in itself. In contrast, our intent is to draw out *ethical* aspects of design, and the roles of technology designs in implicitly and explicitly making 'claims' about the world (Carroll & Rosson, 1992). Thus, while it is not a 'critical design' in the sense of creating a design artefact intended provocation, it is a re-representation of a design as a case, in such a way as to re-represent that case through the lens of reflection-on-action, inviting us to explore ethical texture and imagine alternative designs. To create a *reflective design case*, we developed the following steps, distilled into a model in Figures 1–3.

In the following sections, we apply this model, and through this application exemplify its wider use. We first outline a brief learning situation from which our case arose. The case norms and conventions describe the kinds of socio-material resources of particular

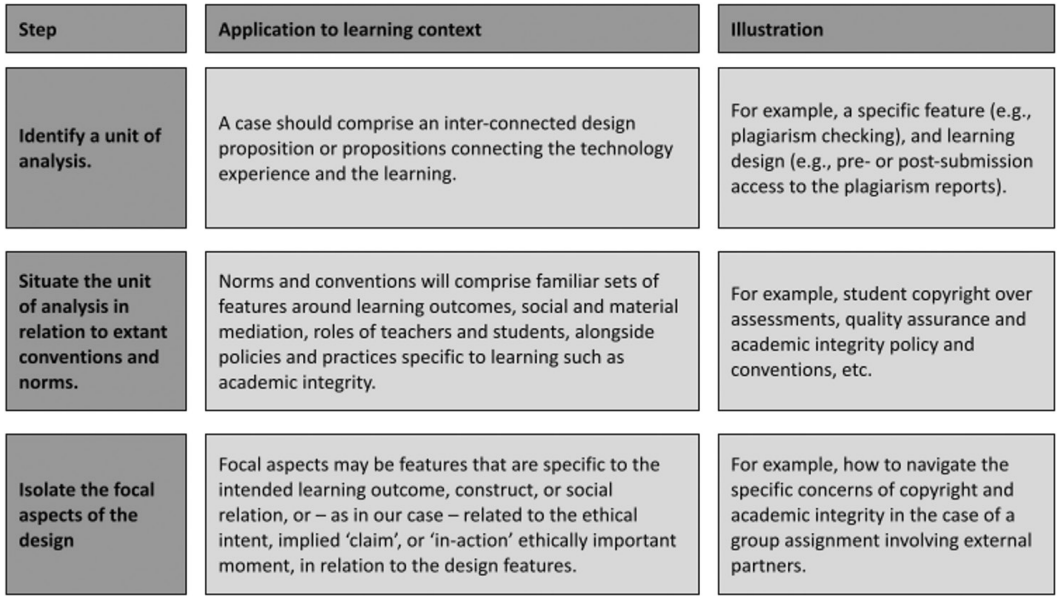

FIGURE 1 A reflective design case model and illustrative example based on analytics of student writing (the authors release this figure under a CC-By licence).

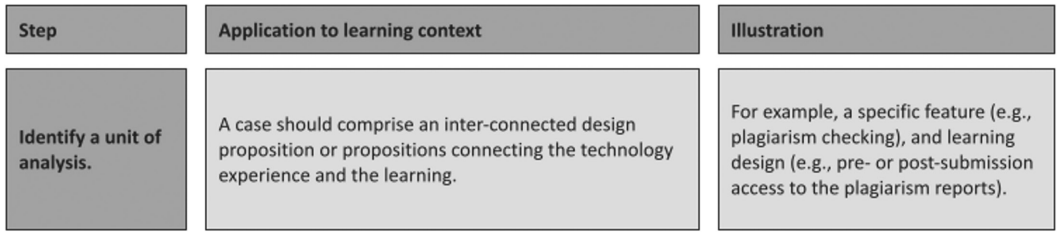

FIGURE 2 Case focus: A reflective design case model and illustrative example based on analytics of student writing (the authors release this figure under a CC-By licence).

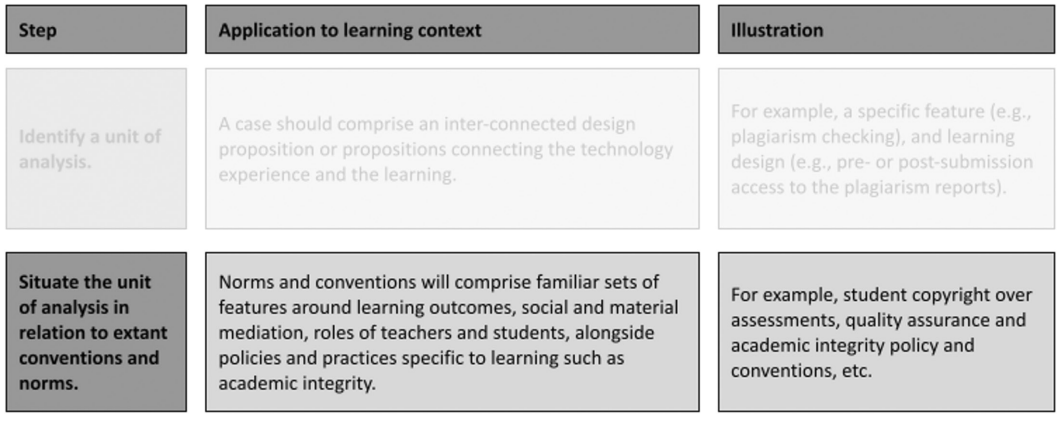

FIGURE 3 Norms and conventions focus: A reflective design case model and illustrative example based on analytics of student writing (the authors release this figure under a CC-By licence).

relevance to our case—the design, development and implementation of a particular tool in an Australian university—while in other cases, local norms and conventions, and those related to disciplinary and pedagogic practice (rather than close-to-practice research) would be emphasised more. The particular ethical focal points are then discussed, drawing on these features to highlight micro-ethics in action.

OUR CASE

Case: Unit of analysis

The unit of analysis for our case is the design, development, and implementation of a writing analytics tool (AcaWriter) at a large metropolitan Australian university (the University of Technology Sydney). Specific features of the design will be drawn out below. The implementation involved a writing analytics tool designed to provide automated feedback on university student academic writing, of a traditional analytical kind (essays, science papers, etc.). The case implementation occurred over approximately a six year period from 2014, with details of the tool available elsewhere (Knight, Buckingham Shum, et al., 2018; Knight, Shibani, et al., 2020),[3] alongside key aspects of the design patterns (Knight, Shibani, et al., 2018).

Case: Norms and conventions

In considering the norms and conventions, the intent is to note key resources that—by convention or by policy—are intended to be drawn on in the context of our case unit. We draw attention to these from most global, to most local context, with respect to their connection to the micro-ethics of our case (summarised in Appendix B).

Institutional context

The 'Common rule' or Belmont principles (Office for Human Research Protections [OHRP], 1978) underpin research ethics, with three core principles: *autonomy* (respect for persons); *beneficence* (maximise benefits and minimise harms); and *justice* (the benefits and risks of research should be distributed fairly). In Australia, Human Research Ethics Committees (HRECs)—equivalent to the Institutional Review Board (IRB) and similar bodies internationally—are charged with oversight of research ethics. As elsewhere, human ethics is legislated, and in Australia the National Health and Medical Research Council (NHMRC) is charged with creating and publishing the "National Statement on Ethical Conduct in Human Research" (National Statement, 2018).

Alongside ethics processes, there is also a raft of guidance and policy around integrity, focusing on issues such as ensuring broad access to research (eg via open access policies), and misconduct (eg via fraudulent data), with an aim that the human endeavour of research should improve our material conditions (OECD, 2007). University research is also bound by a range of relevant national policy such as the Privacy Act, and its instantiation into university Data Governance and Information Security policies. These intersect with the HREC insofar as HREC processes encourage (or, latterly, require) creation of material resources such as a "Research Data Management Plan" (RDMP) outlining the secure collection, storage, deidentification, analysis, sharing and deletion of data; and Impact Canvas, outlining the ways in which the research will benefit society (principles of beneficence, and justice). These tools feed into attempts to promote consideration of research 'soft impacts' (Swierstra, 2015),

side-effects of our actions that are not well captured by risk-focussed HREC models, insofar as they encompass potential consequences for political, economic, social, technological, legal, and environmental (PESTLE) domains,[4] for example in the edtech space, of potential for technologies to impact the labour relations of academics and tutors.

While these institutional bodies are charged with accountability for implementing formal processes and policies, responsibility for ethical conduct of research rests with all researchers. That is, while formal structures exist, it is incorrect to think of these as arbiters of ethical action, and although the phrases, "ethical approval was given", or indeed colloquially, "we got through ethics", are often heard, this is not an accurate representation of the role of the HREC, or the importance of continued ethical reflective practice in the conduct of research and practice (BERA, 2018).

Practice and the resources of practice

Education research is distinct in often involving approaches that are close-to-practice, that is, "any research that focusses on educational practices in order to better understand or improve them." (Wyse et al., 2018, p. 1). The principles of autonomy, beneficence and justice are thus negotiated in this close-to-practice context, in researcher-practitioner dialogue aiming to advance both theory and practice. The translation of these principles into *educational implications* is clear; our research ethics are not just about minimising risks of harms, but about conducting research that supports agency, that develops learning and understanding of that learning, and that addresses inequalities for fairer educational outcomes. This disciplinary and professional context is key to the specific ways in which we navigate beyond principles and structures.

Cases, both from educational and other research, provide some resources to support this navigation, as does public discourse around ethical application of technologies (see for example, the AI, Algorithmic and Automation Incident and Controversy Repository AIAAIC, 2021). However, as also noted by Bezuidenhour and Ratti (2021), these cases tend to (1) be clear in their ethical direction (ie, uncontroversially controversies, or clearly demonstrating maleficence) and (2) coarse-grained (ie, about a particular feature and incident), such as the case of a university president saying: 'drown the bunnies…put a Glock to their heads', in the context of use of predictive analytics to remove students unlikely to complete (Jones et al., 2020). Recent reports involving discussion with various stakeholders regarding AI and its role in learning provide further resource. In the local context, these include a New South Wales Department of Education magazine issue focused on "Teaching and technology: educating in a machine age" and the Australian Federal government 'Artificial Intelligence: Australia's Ethics Framework' artefact and consultation submissions. The former includes discussion with an academic regarding AI in education in NSW (Southgate, 2020) and a piece on AI ethics principles which does relate to two cases (Selinger & Vance, 2020). The latter includes over 130 consultation submissions, However, here too most responses were not education focused (fewer than 10), with the few examples given either clear-cut, or a coarse granularity. Therefore, while extant cases provide some resource, their lack of focus on ethics-in-action and the granular design decisions in design for learning falls into the concerns above regarding a lack of 'texture' in ethics discussions.

Another growing approach in technology design is that of 'value sensitive design' (Friedman, 1996; Friedman et al., 2002), intended to provide a design model that recognises the ways in which our technologies are imbued with values, that impact stakeholder groups in different ways. Briefly, the approach takes a technology, identifies its stakeholders, and benefits and harms for each groups, which are mapped onto values, and analysed for key issues and conflicts in these values, to feed into the design structures. However, while the

approach has received recent attention, becoming part of the conventions in the field (see, eg Chen & Zhu, 2019), it does not offer an adequate model in our case for three reasons: (1) it undertheorizes the technology-use nexus in the context of educational technologies, designed for particular pedagogic contexts or tasks; (2) the analysis of benefits and harms as distinct does not recognise their inseparable nature, and the tradeoffs we make in prioritising one over another; (3) it does not provide conceptual resource for the ethics-in-action or run-time design decisions made in learning contexts, and the kind of close analysis of interest here (for further critique, see Reijers & Gordijn, 2019).

More broadly, then, resources such as the ethics guidelines described above, including those produced by learned societies (such as, BERA, 2018), alongside wider discourse on technologies and research in society, are entwined with activity and discussions as researchers navigate the ethics of our everyday decisions. Additional resources can be found in artefacts such as the 'Ethical Matrix' (Tangen, 2014) and research ethics taxonomy (Beauchemin et al., 2021; Sieber, 2004). The former of these (Ethical Matrix) analyses the Belmont principles with respect to different stakeholder groups: research community; research participants; users of educational services (children/pupils, parents); students in teacher training and special education training; practitioners; and policy makers. The latter (taxonomy of research ethics topics) analyses the principles (and their instantiation into polices such as informed consent) across components of a research programme (Researcher-participant communication; acquisition and use of data; external influences on research; risk and benefit; theory, method and design).

These resources provide helpful framing for the norms and conventions of our work. However, educational research, particularly close-to-practice work, has a set of distinctive features that challenge application of these approaches directly. Therefore, we draw on Beauchemin et al. (2021) and their use of Sieber's (2004) taxonomy of research ethics topics, alongside Tangen's (2014) 'Ethical Matrix', to refocus their analyses (from Tangen's stakeholder concerns across principles; and Beauchemin/Seiber's of principles across different aspects of the research journey) on a typical 'close-to-practice' design approach: the impacts on stakeholders across cyclical research stages of: design, development, evaluation and implementation of a research program.

Case: Focal aspects—Dilemmas? We've had a few…

To discuss the focal aspects of the design case, we use an adapted model that dissects the 'ethically-important-moments' by broad design phase of the research, noting specific features – informed, but not bound by – the range of ethical resources described above. The key tensions are highlighted under questions that arose during critical moments. Specifically, we discuss focal aspects of the design, with respect to stakeholders including: researchers, student-and-staff participants or users of the technology, practitioners (including tutors and academics), at each of the stages, as shown in Figure 4, with summaries of each stage provided in Appendix A.

Inception and design phase: How do we design for impact?

The Academic Writing Analytics project that gave rise to AcaWriter began in 2014, with the aim to develop *automated formative feedback* on *university student writing*. The choice to focus on writing was grounded in (1) the evidence that this is an area students find challenging, and that it is challenging to support, and (2) the alignment of this evidence with the perceived needs at the specific university. It was perceived that new tools, building on natural

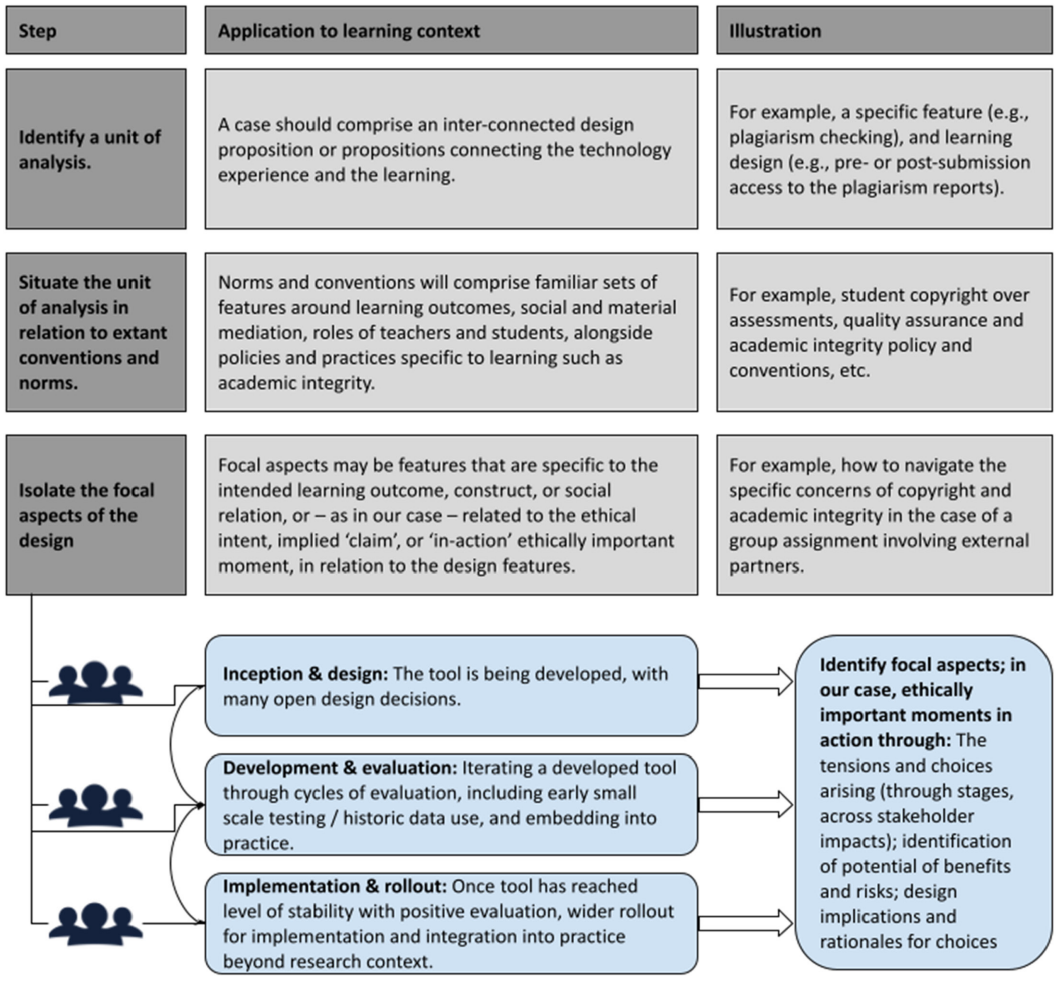

FIGURE 4 A reflective design case model and illustrative example (the authors release this figure under a CC-By licence).

language processing (NLP), provided opportunity to address these concerns, and in particular that an existing parser could be adapted for testing towards this end. Three key issues arose at this stage, characterised around the overarching concern of opportunity costs: If we introduce this change, what other activities are we preventing or downgrading?

What stakes are appropriate?

Our first tension was around the appropriate stakes for automated feedback, as formative, or summative assessment. Our decision was to target formative assessment, on the basis that: (1) developing student writing skills is hard to support, and thus an important target area; (2) the tool's automated feedback was designed for humans to interpret and act on, affording potential for learning from this interpretation; and (3) the tool's NLP was not designed for use in automated assessment (eg it had no knowledge of curriculum content), making this an inappropriate extension. This decision also maintained status quo around labour relations (ie, it did not make human marking of assessments redundant), while augmenting human capacity with machine feedback. We were aware of concerns that: (1) our approach would require human interpretation and thus did not offer the 'quick fix' of more summative

approaches; and (2) the design would give little direct insight to instructors. However, our design decisions were based around balancing of: potential risks (of poor feedback, and low buy in from stakeholders), and benefits (formative feedback being perceived to have higher value in this context); and that this was particularly important in early stages of evaluation, while also being a sustainable model.

On this basis, we worked to design a tool with a student-oriented feedback interface, while working with instructors to ensure alignment of the kinds of feedback provided to their writing tasks and assessment criteria.

What does this imply we value?

In this stage, then, we were also working with instructors to understand where they wished to target feedback. We built on an existing tool, developed to identify 'rhetorical moves' in scholarly literature. We adapted the tool to this learning context, including by understanding the particular language of the disciplines for early adopters (see, Knight, Buckingham Shum, et al., 2018). In this process we sought to ensure that the kinds of rhetorical moves we identified were in fact salient to the disciplinary writing context; that is, that were not implying a focus on features that were not relevant. Of course, a tool could provide feedback on a range of things, including topic coverage, macro-level argument features, and surface features including complexity, sentence/paragraph length, and grammar/spelling. What does it imply about our priorities to focus on one or other of these? And what are the opportunity costs of such a focus? At this point we also had instructor discussion that requested spelling and grammar checking, which we viewed as having potential to over-emphasise low-level feedback (a common issue in writing feedback), and introducing opportunity cost for both development resource, and student focus; however, this issue was discussed more broadly (see next point).

Having developed a student-oriented feedback interface, we nevertheless did not proceed to target students directly, but rather worked with instructors to embed the interface into wider pedagogic context (where the wider context of writing could also be emphasised).

Are there barriers to using the feedback?

As noted above, one feature that was requested was spell checking. Spell checking seems like a banal example of a non-ethical feature, and serves as a good example of micro-ethics in practice. The tool design is based on implementation of a rule-based NLP system that models syntactic relations between concepts. Therefore, users with spelling or grammar errors in their writing are less likely to receive feedback regarding the presence of the identified rhetorical features. This has consequences for provision of equitable feedback, beyond generic instruction that the authors should review their writing. In contrast, though, by including surface level feedback (primarily spelling and grammar), we might overemphasise these features over the rhetorical structures. Moreover, given constrained resources, introducing features presents opportunity costs for areas of other development. An additional reason for excluding surface level feedback was that we were mindful of the need to structure feedback for action, and feedback messages to correct a spelling mistake for example were considered a distraction from higher order features of writing. As a result, early discussions centred on whether we should make all feedback visible in a single interface, or if there should be mechanisms to prioritise feedback to scaffold it being addressed.

While surface level feedback is important, including both for making rhetorical structure visible for readers and for the NLP tool, we decided not to include feedback on these features directly. Instead, we decided to flag the limitations of the tool to users and embedded it in instructional contexts where other tools performing a similar function (such as Grammarly) were made available.

Development and evaluation phase: What is this data for?

Can we use that data for this?

In early phases of development, we sought to test and evaluate the tool's suitability (1) for identifying rhetorical structures in student writing, across disciplines, and (2) the potential for these structures to be related to grade. There are no openly available corpora of student writing, presenting a challenge for such testing. As a result, we used historic institutional data obtained from assignment submissions. Specifically, we obtained grade information and original submissions from our online systems, in negotiation with instructor gatekeepers, based on the institutional privacy policy,[5] and a consent waiver (NS 2.3.10, National Statement, 2018, see details Appendix B).

A tension arose in that although removing names from spreadsheet grade data is simple, student IDs may still be used to easily re-identify data, and—moreover—free-text submissions in the form of document assignments often contain identifiers, spread through headers, footers, coversheets, and the body text itself on occasion; where identifiers have typos in them (eg a number wrong in an ID) they are much harder to remove, while increasing risk of re-identification. Significant work went into identifying and removing these where possible, alongside reporting of only aggregate statistical data, and composites or aggregate exemplars where qualitative samples were used in reporting.

The design was specifically chosen with respect to the national statement (bracketed letters give principles in 2.3.10), to avoid impacts on the *participants' welfare* (*g*), except insofar as the potential for the tools to provide broad benefit, consistent with the privacy policy, in which case results and the tool would be advertised internally. To maximise this likelihood, reduce risk further and increase benefit, we worked with academic stakeholder partners to act as gatekeepers to the student information in their (historic) courses, also facilitating cross-discipline comparisons to be made. We did not consider the analysis we conducted would impact *student ability to derive financial benefits* (*h*) from the work in future, nor was the consent waiver *prohibited under any law* (*i*) that we were aware of.

Nevertheless, while the criteria for a consent waiver were met (and it was approved by the HREC), and we believe the benefits outweighed any risks in this case, we acknowledge the tensions in such waivers not only of direct risks through privacy breaches, but also of risks to trust in the research process.

What stakes are appropriate here?

At the design phase we considered the broad question of the appropriate stakes for the feedback, and how to design a tool to support that. At the development phase though, this must be operationalised into a practical environment for users to engage with the tool. But what stakes are appropriate here, in the specific context of a disciplinary classroom? To put it another way, designing for formative feedback in principle, is different to designing specific opportunities for formative feedback in practice.

At this stage, our key micro-level decision was in how to design a task to expose students to the tool, while acknowledging the unknown benefits of the feedback for learning, and recognising the harms inherent in opportunity costs; that is, that for every task undertaken, students miss out on alternative tasks (and that this might impact teacher evaluations).

To address these concerns, we developed tasks that permitted increasingly higher quality evidence of impact, corresponding with increasing integration into the core learning design (Knight, Shibani, et al., 2018). In the first iteration, students optionally engaged with the tool after they had submitted their assignment—removing the concern it might be a distraction—alongside conducting a self-assessment task. Subsequent tasks, for example, involved reviewing sample texts that were marked up with feedback from the tool, where the text

engagement itself—independent of the markup—was pedagogically useful. Later versions of formative tasks involved students using the tool on their own texts, and indeed questioning the feedback itself again as a pedagogically useful design (Shibani et al., 2022).

How can we evaluate in practice?

In parallel to concerns regarding the appropriate stakes, and the iterative increasing of exposure to the tool, were decisions made around the evaluation of impact. The issue arising is one flagged by Tangen (2014): the trade-off between quality and ethics in educational research. As they note, "Quality of research here refers to both internal criteria such as validity, reliability, and trustworthiness and external criteria such as the relevance and usefulness of research-based knowledge for practice and policy-making." This is a view echoed by Gutiérrez and Penuel (2014), and shared by us, in highlighting that we need not diminish quality in order to foreground relevance to practice as a criterion for rigour. Central to this issue, is that our research should inform practice, and in work like our own, we should not be imposing measures that have no direct benefit to either students or teachers.

However, we recognise the tensions here; indeed, a peer review comment from an article review provided a resource highlighting the issue, noting the deficiencies of typical assignment grading (which we had used) both for inter-rater reliability and measurement validity.[6] Clearly course assignments are used every day, and while they may have limitations, in practical contexts graders engage in practices supportive of reliability (moderation, use of exemplars, etc.), although these do not lend themselves to reliability measurement. As a result, we drew on the data available, and extensions of the existing well established tasks, to develop a suitable evaluation model aligned with practice (Hoadley, 2004).

Development and evaluation phase: What are the power dynamics at the research-practice nexus?

How do we navigate boundaries of researcher, participants, and student?

Education research frequently sits at the research-practice nexus, in learning analytics translating into stakeholder academics acting as both practitioners and researchers, sometimes developing technologies themselves, and sometimes partnering to do so (Shibani et al., 2020). Similarly, classroom research often involves aspects of typical interaction that are not considered a component of the research (the norms of the class, wider sequences of learning tasks, and so on), and thus participants have the dual role of students both in the sense of this wider context, and in that their participation centres around their learning environment.

Here we highlight one feature of this tension in the form of academics as researchers and practitioners, and the norms of research integrity described above. If academics are treated as participants, typically we would seek their consent, treat their data confidentially, and although we might seek confirmation regarding conclusions in qualitative work, we would not anticipate including them as authors. If they are co-researchers, then these things do not hold, but the power dynamic in their relationship with their students then shifts.

In our case, we invited academics on to a subset of papers as co-authors, reflecting their significant contribution to the inception, design, and execution of the work (and to the manuscripts themselves). This in itself introduced a tension; in other work we did treat the academics as participants, and went through an ethics application to conduct interviews with them. However, in that application – and consent form and discussion – we made clear that although we would use pseudonyms in any publication, they would likely be re-identifiable, in part because we had co-authored about our collaborations together.

How do we balance fidelity and flexibility?

While the overarching tool integration was co-designed with the academics, grounded in theory and practice, as in many contexts, delivery was executed by a range of people, in this case tutors. In interviewing the instructors, they noted that "tutor involvement could be a potential factor that affected how the intervention was delivered to students. Tutors are the people who facilitate activities in some classes, and their involvement thus plays a key part in students' engagement" (Shibani et al., 2020, p. 7). Although tutor feedback was largely positive, in part because of the support provided by the instructors, we know from discussion—not formal data capture—that there was variation in tutor implementation, with specific formative tasks augmented by the tool moved in sequence, given variable time in a classroom session, or connected to varying aspects of the intended learning outcomes. This is similar to other contexts where pedagogic strategies and curricula are developed by a range of experts for practitioner use. In these cases, what is the appropriate level of flexibility in execution, given a desire to ensure fidelity to the intended intervention, with respect to tutor autonomy, responsiveness to student needs, and research rigour (see, for example, Gelmez-Burakgazi, 2020)? The concern then is in respecting teacher professionalism, while—in this phase—also seeking to make claims regarding impact grounded in similar 'treatment', and in later phases seeking to implement research backed strategies.

In our case, our aims were to (1) support alignment of research and teaching needs through the task design, and communicate this rationale in collaboration with the instructors in order to maximise tutor buy-in, while also (2) developing our design in such a way that it could be adapted to local contexts, providing for some flexibility (see, Knight, Gibson, & Shibani, 2020, sec. 5; Shibani et al., 2019).

Is this research or 'quality improvement'?

The overarching tension in both cases, and indeed across the research cycle, is around distinctions between research, and improvement as part of ordinary operations. As educators, researchers of learning, and employees of institutions with education as their primary purpose, we have a duty to understand and develop learning processes. As Griffiths (2020) notes, one lens through which to see these questions is in a distinction between the history of research ethics (tracked back to the Nuremberg code), and the history of operations research, which has a similar longevity, but a greater focus on organisational improvement processes and appropriate use of data for legitimate purpose.

In education, these issues play out in learning analytics research, as well as much other educational research, and in use of commercial edtech. In the latter case, as Griffiths (2020) notes, there have been cases of commercial companies changing and testing aspects of products without consent, and discomfort around this. Perhaps in these cases a part of the concern is the commercial component, that companies may make such changes without any consultation, and a wider context of marketisation and centralisation of decision making. Moreover, these issues are contested, and evolving, with no clear demarcation between spaces in either education or health—another field that has grappled with these issues—or indeed other areas.

In our case these issues play out in a number of ways. A standard HREC process expects participants and researchers to be noted separately and in advance, alongside provision of research materials. However, in our work, while principles for materials can be created, many are created by educators 'on the fly', or outside of timelines for a HREC process. Similarly, as noted above, often educators are co-investigators, practitioners, and participants. Exclusion of such research from the body of knowledge is a problematic omission, and thus our strategy has been to have clear understanding of the learning situation into which strategies will be deployed, in order to provide patterns that can be adopted and adapted following a strategy that provides a methodological rigour, and rationale to the HREC.

Implementation and evaluation: Is it ready to use?

How can we scale our evaluation?

In our final set of tensions, the central issues are around how to evaluate technologies beyond opt-in activities, and how to transfer these evaluations to ongoing process of improvement. In education research there is often tension around the use of between groups testing (such as A/B testing, or control group comparison). However, such comparisons provide important data for statistical analysis to evaluate impact. Of concern here are questions around whether students are disadvantaged if they do not receive automated feedback, or/and the opportunity costs of engagement with tasks either with, or in the absence of, some augmentation intervention.

In addressing these concerns our design sought to consider the existing learning context, and the ways in which tasks could be designed that would provide pedagogic value regardless of the tool augmentation. In this way, we sought to ensure no student was disadvantaged by, for example, receiving what might be considered a placebo treatment, or a distractor task. Moreover, even in this phase, the strength of evidence for impact was not strong enough to indicate the tool would provide benefit over the usual task delivery.

What do we mean by scale?

A further concern in implementing is in understanding scale of impact, both in evaluation and subsequent implementation based on the resources provided to academics. At a basic level, impact might be thought of in terms of number of users, or amount of learning (an obviously contested notion) (Knight, Gibson, et al., 2020). This tension was live in institutional context, where there is pressure to ensure funding is delivered on, and "number of users" is an obvious and easy metric for funders to request. This can lead to simply targeting increased use, for example through making the tool easier to use, or targeting larger class cohorts, over increases in learning or targeting users based on justice principles. This can also lead to a focus on models that work in large—but perhaps not transferable—contexts, over implementation approaches that support adoption and adaptation across contexts.

In our work we have sought to develop tools in tandem with resources that support implementation for learning, for example through examples of aligning the tool outputs with assessment criteria, and suggested tasks that are well grounded in evidence such as self-assessment. It is through this task-based approach that we have tried to demonstrate impact, and foster implementation.

When is it ready?

As Wise et al. (2018) ask: "when are learning analytics ready, and what are they ready for?" (p. 1). This question not at face value an obviously ethical concern—has underpinned our approach to research. At what point is a tool mature enough to roll out, independent of researchers? This is a particularly interesting question in learning contexts, in which we aim to build up students' own ability to evaluate work—their evaluative judgement (Boud & Falchikov, 2007) an aim that may be supported just as well, if not better, by 'imperfect' analytics (Kitto et al., 2018). That is, despite the technical context of our work, our evaluation should focus on improvements in learning over algorithmic improvements.

In our context the implication of this is that the tool was 'ready' enough—albeit still receiving evaluation—when researchers, students and academic practitioners were satisfied they could implement it in alignment with their learning aims, and the evaluation of these implementations indicated that the tool fostered supportive learning tasks, and supported desired changes in student writing (Knight, Shibani, et al., 2020).

CONCLUSION

There is increasing attention to ethical concerns, seen in the growing number of guidelines, checklists and policies, and in discourse across scholarly literature, news media, social media and policy responses. While these artefacts shape and are shaped by our development of learning technologies, providing some resource for thinking about and distilling ethical practice, they tend to be coarse in nature, and do not provide practical insight into micro-ethical reasoning-in-action.

This paper has provided a method—the reflective design case—through which micro-analysis of learning contexts and technologies may be conducted. Our case demonstrates the approach and how it provides a novel lens onto design decisions, and their rationale, that are unlikely to be captured through other approaches. A challenge and focus for future research in broader application of approaches such as this, is in understanding the resource costs and learning needs at appropriate points in the design, and to balance the benefit of making cases explicit (as we have done), versus the understanding of these processes as part of the ethics of the everyday for many practitioners. Our particular case provides specific insight into both the kinds of issues that may arise in designing, developing, and implementing learning analytics, and how we have navigated these issues. Of course, these are examples of an illustrative case involving a particular set of people and contexts. However, sharing these, adopting the method we propose, is intended to support the field in developing further nuanced and evolving approaches to practical reasoning in learning analytics. Through developing and researching the application of approaches such as the reflective design case, we aim to augment the tools, guidelines and principles that provide resources for our ethics, with a method to develop reflection and example cases. Through these cases, examples of ethics-in-action can be shared and learnt from to understand how people apply ethical principles in practice, and to foster a textured approach to understanding micro-ethics in practice.

ACKNOWLEDGEMENTS

Our thanks to the participants and collaborators across the research projects within this case. Our thanks also go to the Human Research Ethics Committee (HREC) members and secretariat, who play an important role in the ethical conduct of research, and with whom we have had many productive interactions. Our thanks to Nicole Vincent for discussions particularly regarding soft impacts and technology in society. Thank you to our colleague Tessy Cerratto-Pargman who kindly provided comments on a draft of the manuscript. Open access publishing facilitated by University of Technology Sydney, as part of the Wiley - University of Technology Sydney agreement via the Council of Australian University Librarians.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

The paper is not based on production of original data, and there is thus no data to release.

ETHICS STATEMENT

Human research ethics in Australia is covered by a National Statement on Ethical Conduct in Human Research, as described in detail in this paper. Appendix B sets out the detail of the ethics process under applications: ETH15-0078, ETH16-0285, ETH16-0675, ETH17-1176, ETH18-2263, ETH18-3080, ETH19-3475.

ORCID

Simon Buckingham Shum https://orcid.org/0000-0002-6334-7429

ENDNOTES

[1] At the time of writing, 67 papers refer to "educational technology". Similarly, a simple Google Scholar search for "educational technology" AND "situational ethics", returns only 154 papers in total.

[2] While we are aware of the discussion around reflexivity and reflection, here we treat them as broadly synonymous, using reflective except where authors specifically refer to reflexive.

[3] Academic Writing Analytics project, University of Technology Sydney: https://cic.uts.edu.au/tools/awa

[4] We note at time of writing (October 2021) Google Scholar returns 28 results for: "educational technology" AND ("PESTLE" OR "STEEPLE") AND [each of the PESTLE elements]. (i.e. articles explicitly using that model), including an interesting discussion of tradeoffs in growth of online education in higher education (Morris et al., 2020).

[5] On first enrolling and logging into IT systems, students agree to a privacy statement that in 2015 included the line: "information, including data automatically generated through my access to and use of [the LMS], may also be used to provide me support and may be assessed so that additional services and facilities can be offered to me to support my studies. Analysis of deidentified information may also be used to inform UTS quality improvement initiatives." For a detailed discussion of these statements, informed consent, and the nature of operational vs research ethics, see Griffiths (2020).

[6] Having reviewed the publication policies and Committee on Publication Ethics (COPE) guidance, we believe that while it is acceptable for us to share this example, a strict interpretation of copyright is typically applied to reviewer reports, hence we refrain from direct quote here.

REFERENCES

AIAAIC. (2021). *AIAAIC - AI, algorithmic and automation incident and controversy repository*. https://docs.google.com/spreadsheets/u/0/d/1Bn55B4xz21-_Rgdr8BBb2lt0n_4rzLGxFADMlVW0PYI/htmlview#

Bardzell, J., Bardzell, S., & Stolterman, E. (2014). Reading critical designs: Supporting reasoned interpretations of critical design. In M. Jones, P. Palanque, A. Schmidt, & T. Grossman (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing System* (pp. 1951–1960). ACM. https://doi.org/10.1145/2556288.2557137

Beauchemin, É., Côté, L. P., Drolet, M.-J., & Williams-Jones, B. (2021). Conceptualising ethical issues in the conduct of research: Results from a critical and systematic literature review. *Journal of Academic Ethics*, *20*, 335–358. https://doi.org/10.1007/s10805-021-09411-7

BERA. (2018). *Ethical guidelines for educational research, fourth edition (2018)*. BERA. https://www.bera.ac.uk/publication/ethical-guidelines-for-educational-research-2018

Bezuidenhout, L., & Ratti, E. (2021). What does it mean to embed ethics in data science? An integrative approach based on microethics and virtues. *Ai & Society*, *36*(3), 939–953. https://doi.org/10.1007/s00146-020-01112-w

Boud, D., & Falchikov, N. (2007). *Rethinking assessment in higher education: Learning for the longer term*. Routledge. https://books.google.co.uk/books?hl=en&lr=&id=GJt9AgAAQBAJ&oi=fnd&pg=PP1&ots=GouS42kgtU&sig=Ld6-4lSwA7H1XpNWHmmDkY0u_Es

Carroll, J. M., & Rosson, M. B. (1992). Getting around the task-artifact cycle: How to make claims and design by scenario. *ACM Transactions on Information Systems (TOIS)*, *10*(2), 181–212. https://doi.org/10.1145/146802.146834

Chen, B., & Zhu, H. (2019). Towards value-sensitive learning analytics design. In S. Hsiao, J. Cunningham, G. Lynch, C. Brooks, R. Ferguson, & U. Hoppe (Eds.), *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (pp. 343–352). Association for computing machinery. https://doi.org/10.1145/3303772.3303798

Frauenberger, C., Rauhala, M., & Fitzpatrick, G. (2017). In-action ethics. *Interacting with Computers*, *29*(2), 220–236. https://doi.org/10.1093/iwc/iww024

Friedman, B. (1996). Value-sensitive design. *Interactions*, *3*(6), 16–23. https://doi.org/10.1145/242485.242493

Friedman, B., Kahn, P., & Borning, A. (2002). *Value sensitive design: Theory and methods* (Technical Report 02-12). University of Washington. http://www.urbansim.org/pub/Research/ResearchPapers/vsd-theory-methods-tr.pdf

Gelmez-Burakgazi, S. (2020). Curriculum adaptation and fidelity: A qualitative study on elementary teachers' classroom practices. *Issues in Educational Research*, *30*(3), 920–942. http://search.informit.org/doi/abs/10.3316/INFORMIT.465246677972448

Gray, C. M., & Boling, E. (2016). Inscribing ethics and values in designs for learning: A problematic. *Educational Technology Research and Development*, *64*(5), 969–1001. https://doi.org/10.1007/s11423-016-9478-x

Griffiths, D. (2020). The ethical issues of learning analytics in their historical context. In D. Burgos (Ed.), *Radical solutions and Open Science* (55). Springer; Scopus. https://doi.org/10.1007/978-981-15-4276-3_3

Guillemin, M., & Gillam, L. (2004). Ethics, reflexivity, and "ethically important moments" in research. *Qualitative Inquiry*, *10*(2), 261–280. https://doi.org/10.1177/1077800403262360

Gutiérrez, K. D., & Penuel, W. R. (2014). Relevance to practice as a criterion for rigor. *Educational Researcher*, *43*(1), 19–23. https://doi.org/10.3102/0013189X13520289

Hoadley, C. M. (2004). Methodological alignment in design-based research. *Educational Psychologist*, *39*(4), 203–212. https://doi.org/10.1207/s15326985ep3904_2

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2021). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, *32*, 504–526. https://doi.org/10.1007/s40593-021-00239-1

Johansen, M. B., & Frederiksen, J. T. (2021). Ethically important moments—A pragmatic-dualist research ethics. *Journal of Academic Ethics*, *19*(2), 279–289. https://doi.org/10.1007/s10805-020-09377-y

Johnson, R. L., Liu, J., & Burgess, Y. (2017). A model for making decisions about ethical dilemmas in student assessment. *Journal of Moral Education*, *46*(2), 212–229. https://doi.org/10.1080/03057240.2017.1313725

Jones, K. M. L., Rubel, A., & LeClere, E. (2020). A matter of trust: Higher education institutions as information fiduciaries in an age of educational data mining and learning analytics. *Journal of the Association for Information Science and Technology*, *71*(10), 1227–1241. https://doi.org/10.1002/asi.24327

Kitto, K., & Knight, S. (2019). Practical ethics for building learning analytics. *British Journal of Educational Technology*, *50*(6), 2855–2870. https://doi.org/10.1111/bjet.12868

Kitto, K., Buckingham Shum, S., & Gibson, A. (2018). Embracing imperfection in learning analytics. In A. Pardo, K. Bartimore-Aufflick, G. Lynch, S. Buckingham Shum, R. Ferguson, A. Merceron, & X. Ochoa (Eds.), *Proceedings of the 8th International Conference on Learning Analytics and Knowledge* (pp. 451–460). Association for Computing Machinery. https://doi.org/10.1145/3170358.3170413

Knight, S., Buckingham Shum, S., Ryan, P., Sándor, Á., & Wang, X. (2018). Academic writing analytics for civil law: Participatory design through academic and student engagement. *International Journal of Artificial Intelligence in Education*, *28*(1), 1–28. https://doi.org/10.1007/s40593-016-0121-0

Knight, S., Gibson, A., & Shibani, A. (2020). Implementing learning analytics for learning impact: Taking tools to task. *Internet and Higher Education*, *45*, 100729. https://doi.org/10.1016/j.iheduc.2020.100729

Knight, S., Shibani, A., Abel, S., Gibson, A., Ryan, P., Sutton, N., Wight, R., Lucas, C., Sándor, Á., Kitto, K., Liu, M., Mogarkar, R., & Buckingham Shum, S. (2020). AcaWriter: A learning analytics tool for formative feedback on academic writing. *Journal of Writing Research*, *12*(1), 141–186. https://doi.org/10.17239/jowr-2020.12.01.06

Knight, S., Shibani, A., & Buckingham Shum, S. (2018). Augmenting formative writing assessment with learning analytics: A design abstraction approach. In J. Kay & R. Luckin (Eds.), *13th international conference of the learning sciences: Rethinking learning in the digital age. Making the learning sciences count* (Vol. 3, pp. 1783–1790). International Society of the Learning Sciences. https://doi.org/10.22318/cscl2018.1783

Lu, Q., Zhu, L., Xu, X., Whittle, J., Zowghi, D., & Jacquet, A. (2022). *Responsible AI pattern catalogue: A multivocal literature review* (arXiv:2209.04963). arXiv. https://doi.org/10.48550/arXiv.2209.04963

Morris, N. P., Ivancheva, M., Coop, T., Mogliacci, R., & Swinnerton, B. (2020). Negotiating growth of online education in higher education. *International Journal of Educational Technology in Higher Education*, *17*(1), 48. https://doi.org/10.1186/s41239-020-00227-w

National Statement. (2018). *National Statement on ethical conduct in human research*. National Health and Medical Research Council, the Australian Research Council and universities Australia. https://www.nhmrc.gov.au/about-us/publications/national-statement-ethical-conduct-human-research-2007-updated-2018

OECD. (2007). *OECD principles and guidelines for access to research data from public funding (88180 2007)*. OECD. https://www.oecd.org/sti/inno/38500813.pdf

Office for Human Research Protections (OHRP). (1978). Belmont report: Ethical principles and guidelines for the protection of human subjects of research. https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html

Reijers, W., & Gordijn, B. (2019). Moving from value sensitive design to virtuous practice design. *Journal of Information, Communication and Ethics in Society*, *17*(2), 196–209. https://doi.org/10.1108/JICES-10-2018-0080

Schön, D. A. (1992). *The reflective practitioner: How professionals think in action*. Routledge. https://doi.org/10.4324/9781315237473

Selinger, E., & Vance, A. (2020). Teaching privacy and ethical guardrails for the AI imperative in education. *Future Edge, State of NSW* (*Department of Education*), *Teaching and Technology: Educating in a Machine Age* (3), 30–53. https://education.nsw.gov.au/content/dam/main-education/teaching-and-learning/education-for-a-changing-world/media/documents/Future_EDge_Issue_3.pdf

Sengers, P., Boehner, K., David, S., & Kaye, J. J. (2005). Reflective design. In O. Bertelsen, N. O. Bouvin, P. Krogh, & M. Kyng (Eds.), *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility* (pp. 49–58). Association for Computing Machinery. https://doi.org/10.1145/1094562.1094569

Shibani, A., Knight, S., & Buckingham Shum, S. (2020). Educator perspectives on learning analytics in classroom practice. *Internet and Higher Education*, *46*, 100730. https://doi.org/10.1016/j.iheduc.2020.100730

Shibani, A., Knight, S., & Buckingham Shum, S. (2022). Questioning learning analytics? Cultivating critical engagement as student automated feedback literacy. In A. F. Wise, R. Martinez-Maldonado, & I. Hilliger

(Eds.), *12th International Learning Analytics and Knowledge Conference* (pp. 326–335). ACM. https://doi.org/10.1145/3506860.3506912

Shibani, A., Knight, S., & Shum, S. B. (2019). Contextualizable learning analytics design: A generic model and writing analytics evaluations. In S. Hsiao, J. Cunningham, K. McCarthy, G. Lynch, C. Brooks, R. Ferguson, & U. Hoppe (Eds.), *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (pp. 210–219). ACM. https://doi.org/10.1145/3303772.3303785

Sieber, J. E. (2004). Empirical research on research ethics. *Ethics & Behavior*, *14*(4), 397–412. https://doi.org/10.1207/s15327019eb1404_9

Southgate, E. (2020). Advancing learning through AI: Insights from a NSW teacher-educator and emerging technology researcher. *Future Edge, State of NSW* (*Department of Education*), *Teaching and Technology: Educating in a Machine Age* (3), 12–29. https://education.nsw.gov.au/content/dam/main-education/teaching-and-learning/education-for-a-changing-world/media/documents/Future_EDge_Issue_3.pdf

Stake, R. E. (2003). Case studies. In N. Denzin & Y. Lincoln (Eds.), *Strategies of qualitative inquiry* (pp. 134–164). Sage.

Swierstra, T. (2015). Identifying the normative challenges posed by technology's 'soft' impacts. *Etikk i Praksis - Nordic Journal of Applied Ethics*, *1*, 5–20. https://doi.org/10.5324/EIP.V9I1.1838

Tangen, R. (2014). Balancing ethics and quality in educational research—The ethical matrix method. *Scandinavian Journal of Educational Research*, *58*(6), 678–694. https://doi.org/10.1080/00313831.2013.821089

Wise, A. F., Knight, S., & Ochoa, X. (2018). When are learning analytics ready and what are they ready for. *Journal of Learning Analytics*, *5*(3), 1–4. https://doi.org/10.18608/JLA.2018.53.1E

Wyse, D., Brown, C., Oliver, S., & Poblete, X. (2018). *Close-to-practice educational research*. BERA. https://www.bera.ac.uk/wp-content/uploads/2018/11/BERA-Close-to-Practice_statement_Nov2018-1.pdf

Yeaman, A. R., Eastmond, Jr., J. N., & Napper, V. S. (2013). Professional ethics and educational technology. In A. Januszewski & M. Molenda (Eds.), *Educational technology* (pp. 295–338). Routledge.

How to cite this article: Knight, S., Shibani, A., & Buckingham Shum, S. (2023). A reflective design case of practical micro-ethics in learning analytics. *British Journal of Educational Technology*, *00*, 1–21. https://doi.org/10.1111/bjet.13323

APPENDIX A: OVERVIEW OF ETHICALLY IMPORTANT MOMENTS

Inception and design phase	
Tension	**Design implications**
What are the stakes?	Student (not teacher) oriented tool, for *formative feedback on* writing
What role should automated tools play in formative or summative assessment?	Feedback to augment existing assessment structure
What values does this imply? If we give feedback on this, what does that suggest about what it's important to know?	Early on, we decided to include features to make clear "the machine does not understand your writing" and not to rollout to students without instructor involvement
Who will use this, and how? Are there barriers to the effective use of the tool and its feedback that might result in, or maintain, inequities?	Choice not to include grammar and spelling checking to avoid (1) resource cost, and (2) overemphasising surface level feedback; despite its potential use for some (and its relevance to instructors)

APPENDIX A (Continued)

Development and evaluation phase: Operational and scholarly research	
Tension	**Design decision**
Can we use that data for this? Is consent to a privacy policy clear enough that data may be used for learning research for 'quality improvement'?	Submission of a consent waiver HREC with instructor gatekeepers to the data, and processes to (1) de-identify the data for processing and publication, and (2) recognise that data may still contain identifiers
What stakes are appropriate here? How do we balance opportunity costs around formative assessment designs where we want to minimise risk of harm while maintaining potential benefits?	We developed an iterative learning design that allowed us to increase the integration of the tool as an augmentation of well-grounded pedagogic tasks, alongside increasing evidence supporting potential impact of the tool
How can we evaluate in practice? Researchers typically seek to maximise internal validity and interrater reliability, but this may not align with existing classroom practice or the assessment needs of the students and teachers. How do we balance these needs, and what weight should relevance to practice as a criterion for rigour receive?	To support implementation in an existing context, and to provide for sustainability of the approach through such alignment, we drew on the data available, and extensions of the existing well-established tasks, to develop a suitable evaluation model
Development and evaluation phase: Research practice nexus	
Tension	**Design decision**
How do we navigate boundaries of researcher, participant, and student? Working outside lab contexts often involves participation with people who take on multiple roles across the research, with power dynamics both between researcher-participant, and participant groups	Research evaluation was designed to support learning, with tasks specifically for research requiring opt-in; practitioner participants were invited as co-researchers where appropriate, and consented and engaged as participants by mutual agreement where appropriate
How do we balance fidelity and flexibility? While we co-designed researcher with academic leaders, often teaching is not conducted by those who develop the resources, but by tutors who have some flexibility in delivery. This may impact the claims that can be made regarding interventions, by reducing fidelity	Research was designed to support implementation, while allowing some flexibility in delivery. Tutors were engaged with the work to ensure they understood its aims and potential, but ultimately have some freedom in delivery
Is this research or quality improvement? Human participant research requires particular things, while quality improvement/operations research processes follow different procedures	While the work contributed to quality improvement processes, we ensured participants received appropriate information, and opt-in consent where specific research tasks (such as being video recorded) were undertaken. We did not collect personal sensitive information, and there were clear lines between the data shared across the researchers and academic partners to ensure no unintended impact on learning
Implementation and evaluation: Is it ready to use?	
Tension	**Design decision**
How can we scale our evaluation? As tool evaluation develops, at what point is between-groups comparison inappropriate, and at what point might stakeholders consider embedding tools as a key part of their teaching practice? There are opportunity costs and issues of autonomy that should be thought of separately in research and teaching contexts, although they intersect	Our evaluation model was designed to build over time, augmenting existing practices to support these in class, and to obtain evidence from the range of stakeholders

(Continues)

APPENDIX A (Continued)

Implementation and evaluation: Is it ready to use?	
Tension	**Design decision**
What do we mean by scale? As evaluation develops, how should we balance desire to achieve scaled impact in terms of numbers of users or impact on learning?	While our evaluations were conducted in relatively large classes, we have not sought to ground evaluation in 'numbers of users' instead focusing on alignment of learning impact with intended learning outcomes
When is it ready? At what point are learning analytics 'ready' to be implemented independent of researchers, and outside of specific research study contexts?	Our evaluation focused on learning impact and implementation over algorithmic development, aiming to create resources that would support the tool being implemented into practice while fostering other well-supported formative tasks

APPENDIX B: INSTANTIATION OF ETHICAL FOCI THROUGH NORMS AND POLICY REFLECTED IN RESOURCE OF HREC MATERIAL

The ethics process undertaken in this work is briefly outlined below, following the procedures for HREC applications in the Australian context, at the authors' institution. In addition, a number of other projects were undertaken including on reflective writing, and the application of the discussed tool to higher degree research (HDR) writing; these are not included here.

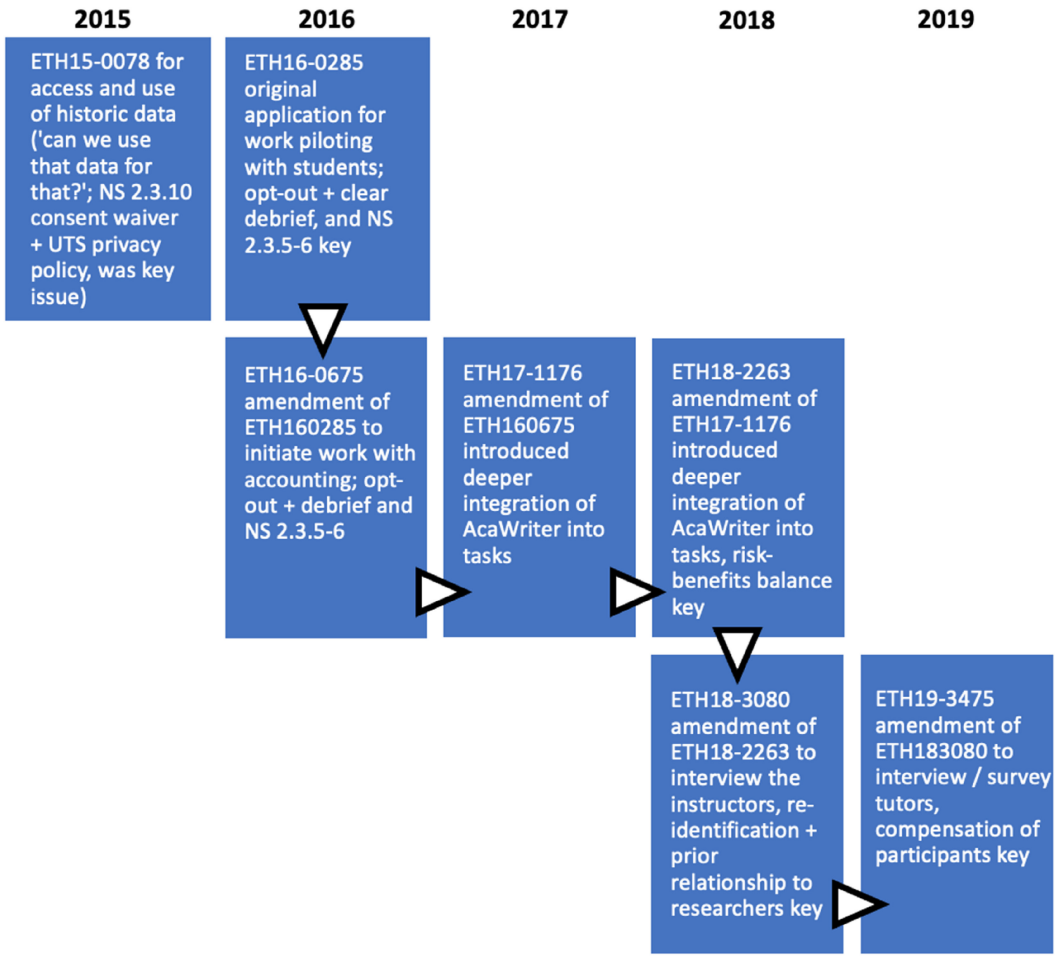

Completing such applications involves bringing together the material resources embedded in that ethics process itself (the HREC form, the National Statement principles, etc.), the specifics of the study to be undertaken (materials, participants, research expertise, etc.), and the negotiation of this between researchers, participants and gatekeepers, and the HREC committee (typically at arms length).

Consent waiver details: A consent waiver was sought on the basis that (referring to specific NS conditions by letter): The research carried *negligible risk* (*a*) since the work had already been assessed and the analysis was conducted on de-identified data such that it would not be possible to reconnect it to the original students. We believe the benefits justified risks from *not seeking consent* (*b*) because the analysis was conducted to understand how to provide formative feedback to students on their writing, and it would be *challenging to obtain consent* (*c*) given that the data was historic, and accessible (as documents of record) via an institutional IT system the students may not have retained access to. We also believe it was unlikely the participants would have *declined consent* (*d*), based on the low risk of the analysis and its potential benefits, and the students' agreement to our privacy policy, which provides broad legitimate purpose for data to be used "so that additional services and facilities can be offered to me to support my studies. Analysis of deidentified information may also be used to inform UTS quality improvement initiatives"[5]. Based on this, we deemed that access to de-identified data should be possible where this would have no impact on student outcomes, but would contribute to 'quality improvement initiatives'. We also sought to ensure *protection of privacy* (*e*) and *confidentiality of data* (*f*) through de-identification of the data at the earliest opportunity (consistent with the privacy policy), and storage on secure machines with access only to the research team.

Consent opt-out details: Under 2.3.5 of the national statement an opt-out approach may be used where it is feasible to contact participants, but the scale and significance of the project means using explicit consent is neither practical nor feasible. Per 2.3.6 the HREC was satisfied that: there was low risk (a), there was public interest (b), the research would be compromised if participation was not near complete and explicit consent would compromise this (c), that reasonable attempts were made to communicate about the research to participants (d), and that they would have opportunity to withdraw over a reasonable period (e), with opportunity to ask questions and decline (f) and maintenance of data security (g) and governance (h), with no specific prohibition (i).

Opt-out was deemed appropriate for the context, given that students were engaging in the learning tasks as part of their standard delivery; augmenting this task, or requesting an additional activity to provide consent from the cohort was not practical, while by this phase of the research our analysis would have been compromised by low participation. Of course, non-assessment items in courses are not typically compulsory, and thus participation in the tasks themselves was not mandated. Students were provided with debrief information including rationale and information regarding opt-out following the target tasks.