Between Privacy and Usability: User Challenges in a Secure Educational Data Enclave

Ryan S. Baker*
Adelaide University
Adelaide, South Australia, Australia
Penn Center for Learning Analytics
University of Pennsylvania
Philadelphia, Pennsylvania, USA
rybaker@upenn.edu

Camille Giordano
Penn Center for Learning Analytics
University of Pennsylvania
Philadelphia, Pennsylvania, USA
camille.giordano.upenn@gmail.com

Debshila Basu Mallick
Rice University
Houston, Texas, USA
debshila@rice.edu

Katherine Fletcher
Rice University
Houston, Texas, USA
kef@rice.edu

Amanda Barany
University of Pennsylvania
Philadelphia, Pennsylvania, USA
amanda.barany@gmail.com

Kirk P Vanacore
Cornell University
Ithaca, New York, USA
kirk.vanacore@gmail.com

Stephen Hutt
Department of Educational Psychology
University of Minnesota
Minneapolis, Minnesota, USA
shutt@umn.edu

Abstract

Concerns about student privacy increasingly limit direct access to large-scale educational data, prompting interest in data enclaves—secure environments that enable in-situ analysis without releasing raw data. While data enclaves are widely used in other domains, their adoption in educational research has been limited. This paper examines user perspectives on MORF (the MOOC Replication Framework), a past attempt to create a publicly available, privacy-preserving data enclave for education data. We report results from a survey of MORF users. Survey responses indicate that MORF enabled research that would otherwise have been infeasible: half of respondents produced publishable outputs using MORF data, and most reported no viable alternative datasets. At the same time, users identified substantial barriers to adoption, including steep learning curves, limited documentation, challenges with data discovery, constrained debugging workflows, and opaque error reporting (many of which stemmed from privacy-driven design choices). User accounts match a prior post-mortem by developers in some areas but not others: while users agreed with developers that privacy protections disrupted iterative research workflows, they were less negative about containerisation and more negative about data documentation than the developers. We conclude by identifying important considerations for future educational data enclaves, emphasising the need for stronger documentation and better privacy-preserving debugging support.

CCS Concepts

• **Applied computing**; • **Education**; • **E-learning**; • **Security and privacy**; • **Privacy-preserving protocols**; • **Human-centered computing**; • **Human computer interaction (HCI)**; • **Empirical studies in HCI**;

Keywords

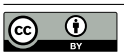
Data enclave, privacy protection, MORF

ACM Reference Format:
Ryan S. Baker, Camille Giordano, Debshila Basu Mallick, Katherine Fletcher, Amanda Barany, Kirk P Vanacore, and Stephen Hutt. 2026. Between Privacy and Usability: User Challenges in a Secure Educational Data Enclave. In *Proceedings of the 38th Australasian Conference on Human-Computer Interaction (OzCHI '26), November 21–25, 2026, Adelaide, SA, Australia.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3847401.3857144

1 Introduction

Concerns about student data privacy remain prominent in discussions about the potential use of large-scale educational data to support basic research and improve learning experiences [11, 25]. While direct access to data is preferred by many data scientists, some data is considered too sensitive (by its owners or by legislation) to share, even if a good-faith attempt is made to remove personally identifying information (PII; [15]). In addition, some data sets are too large for humans to fully deidentify, humans make errors in deidentification [31], and even state-of-the-art deidentification tools do not achieve 100% recall for PII [2, 9, 31]. Other privacy-preserving techniques, such as differential privacy, offer

*Corresponding author



This work is licensed under a Creative Commons Attribution 4.0 International License.
OzCHI '26, Adelaide, SA, Australia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2724-5/2026/11
https://doi.org/10.1145/3847401.3857144

some level of protection but introduce trade-offs in terms of analytic utility and validity [28]. Yet other approaches attempt to create secure environments where research can be conducted without exporting data [13]. These approaches, however, consider privacy in primarily technical terms, leaving fundamentally human-centered questions unaddressed: how do the measures introduced to protect privacy reshape and constrain the way researchers work, and can the systems built to enforce these protections actually be used by the researchers who depend on them.

We ask these questions in the context of researchers' use of data from MOOCs, Massive Open Online Courses. Protecting student privacy is a legal and normative obligation, but it is not completely clear which obligations apply to MOOC data. In the United States, where the platform we will discuss was hosted, the Family Educational Rights and Privacy Act (FERPA) protects personally identifiable information in the education records of enrolled students [6], but its application to MOOCs is contested: the platforms are not themselves FERPA-covered institutions, and open-course participants are frequently not deemed "students" of the offering institution — leading the U.S. Department of Education's own chief privacy officer to observe that MOOC data is seldom FERPA-protected [30, 32]. Australia has no education-specific equivalent; such data is instead covered under the general Privacy Act 1988 (Cth) and its Australian Privacy Principles [23], with state and territory regulations and NHMRC research-ethics oversight also relevant — though whether these apply to a given dataset turns on whether individuals remain reasonably identifiable and on the type of entity (e.g., public or private) holding the data. Compounding this complexity, Australia's regulatory environment is tightening: the Privacy and Other Legislation Amendment Act 2024 introduced a statutory tort for serious invasions of privacy and an explicit requirement for "technical and organisational measures" to safeguard personal information [24]. Given these overlapping and evolving requirements, infrastructure that never releases raw data may prove particularly useful precisely because it can support research and open science practices while protecting participants and limiting institutional risk.

These challenges have prompted calls for the use of data enclaves in education—secure, access-controlled environments that allow researchers to analyse sensitive or confidential data [13]. Such enclaves typically support in-situ analysis tools while limiting direct access to, or export of, raw data in order to protect confidentiality [21]. Managing who may see and use sensitive data, and under what controls, is itself an active concern in human-computer interaction [3], and data enclaves are currently used across multiple domains, including healthcare [17, 22], economics [18], and geospatial research [12].

One of the more significant attempts at creating a publicly accessible enclave for educational data was MORF, the MOOC Replication Framework [10]. MORF was a joint project of three U.S. institutions: the University of Pennsylvania, the University of Michigan, and the University of Denver. At its peak, the platform housed data from millions of learners enrolled in hundreds of MOOCs and supported a variety of research studies that would have been difficult to carry out without access to such specialised data. Nevertheless, despite the distinctiveness of its data and the research opportunities it enabled, MORF saw limited adoption beyond the institutions that developed it [1]. In addition, despite the release of MORF's source code, that code has not been used by other researchers or developers, to the best of our knowledge.

Baker and Hutt [1] offered a perspective, as two of the platform's developers, on the factors that may have contributed to limited adoption. They noted that MORF consistently prioritised participant privacy, but that this emphasis created a persistent trade-off in which usability and (to a lesser extent) functionality suffered, ultimately limiting adoption. They also discussed prospective users' challenges with the platform's core workflow—particularly the need to work through Docker, to plan and implement analyses without being able to directly view the underlying data, and to debug code with limited and delayed feedback—frictions that interrupted the normal iterative cycle of analysis. They further argued that restrictions on allowable outputs, while necessary for privacy, constrained what researchers could do and added additional overhead to the process of conducting analyses, particularly hampering debugging. As a result, in their account, initial enthusiasm often turned into frustration and disengagement, and successful projects generally required substantial involvement from researchers at Penn or Michigan rather than being conducted independently by external teams.

Developers are not always well-positioned to judge users' experiences, however: they may not recognise which frictions matter most, and they may be overly optimistic — or, in this case, pessimistic — about how a system is actually experienced in use (cf. [19]). This gap between designers' and users' perceptions is a longstanding concern in human-computer interaction [7] and is especially consequential for research infrastructure, where the users are themselves researchers whose work is enabled or constrained by design choices made in the name of privacy [14, 27]. Such infrastructure sits at the boundary between two goals that are frequently in tension — protecting participants and enabling research — and the choices made at that boundary shape both how well a platform works and who is able to use it at all. To understand how this tension is experienced in practice, we surveyed former users of MORF, with the platform's developers facilitating access to past users, about the platform's benefits and flaws, drawing on their perspectives to identify considerations for the design of future privacy-preserving research infrastructure, in education and beyond.

2 methods

2.1 Participants and Recruitment

A colleague at the platform team compiled a list of 17 former MORF users who had used version 2.0 or later [e.g. 10], who had never been members of the core development team, to contact via email to take the survey. When multiple members of the same research team all used MORF for the same research project, only one member (the primary contact person) was contacted. The MORF survey was administered online, with the platform team contacting the participants on our behalf to protect participant privacy, and was conducted with the approval of the REDACTED IRB, with all participants consenting.

Thirteen former MORF users responded to the email query and completed the survey, a 76.4% response rate (one participant did not answer every item; their responses are included only where

present). Just over half of respondents (53.8%) were external users who were not affiliated with the institutions that developed MORF (University of Michigan, University of Pennsylvania, University of Denver), indicating that MORF was used by researchers beyond the core development team's institutions. All of the participants were students and/or professional researchers. The majority of respondents (69%) reported having a Master's degree or lower, suggesting that many users were graduate students, research staff, or early-career researchers rather than faculty.

2.2 Survey Composition

The survey[1] was composed of 11 short-answer questions, 20 Likert-scale items, and 4 multiple-choice questions. The questions were developed to ascertain what users were able to accomplish with MORF, what worked well, what didn't work well, and what improvements can be made for future data enclave platforms. Likert questions addressed the usability, clarity of documentation, ease of completing work using MORF, and satisfaction with access to data. Six-part Likert Scale questions were used in order to avoid neutral responses and because of evidence for higher reliability than five-part Likert Scale questions [5]. Short-answer questions prompted respondents to describe in their own words what worked well, what was challenging, and what they would like to see improved in MORF and similar data enclaves. Multiple-choice questions addressed users' comfort level, ease of use, efficiency, and the user-friendliness of the platform. Participants were not required to respond to all questions, and many participants responded to only a subset of questions (more specifically, 6 participants did not answer all questions, and 15 questions were not answered by all participants). For each item, percentages reported below are calculated only out of the participants who responded to that item.

2.3 Analysis

Likert and multiple-choice responses were summarised descriptively. Short-answer responses were analysed qualitatively: two of the authors reviewed the responses and, through discussion, grouped them into recurring categories reflecting what users valued, found difficult, or wished to see improved; we report the proportion of respondents whose comments fell into each category. We conducted interpretive categorisation, as opposed to a deductive qualitative coding with formal inter-rater reliability statistics, because of the small sample size and the lack of an existent qualitative codebook, and because the responses were brief and largely unambiguous. Our aim was to characterise the range of user experiences rather than to establish a replicable coding scheme [16]. This methodological choice is characteristic in well over half of qualitative research conducted in the HCI community [16].

3 Results

We organise our findings around three areas: (1) how researchers used MORF, (2) the strengths in the platform that they valued, and (3) the challenges they encountered. A recurring pattern, which we take up in the discussion, is that the friction users experienced arose less from the privacy mechanism itself than from the surrounding documentation and tooling.

3.1 Uses of MORF

The users reported using MORF for a wide range of research topics, including bias and fairness (17%), tool development (17%), and dropout prediction (17%), student behaviour (17%), to test machine learning algorithms (8%), for "teaching & experimentation" (8%), to understand data management (8%), to study academic work (8%), and for "learning analytics" (8%).

50% of respondents indicated that they had successfully published at least one paper, poster, or other research output that relied on MORF data. This suggests that, for a subset of users, MORF was instrumental in producing tangible scholarly work.

3.2 Perceived Strengths of MORF

Overall, survey respondents reported that they liked having access to the large and unique MOOC database, something that they previously could not access prior to the creation of MORF. One user wrote "The access to MORF database was a huge help", while another said "I really liked the fundamental idea. Being able to have access to datasets folks may not otherwise have access to, with build-in functionality (although limited) to carry out analyses." 41.7% of users saw access to a large dataset as a key benefit of MORF, stating that its benefits included "Easy-ish access to large dataset", "access to historical student data", and "Support multiple researchers for their research projects; streamline data retrieval processes".

The use of Docker was seen as a strength by some users, with one user stating "Docker is easy to debug so it was a good developer experience", and another user noting "I think the best part was the dockerized portion that allows any language to be used" (though this second user was generally negative about MORF otherwise). The code examples were also seen as a plus by some users, with one user noting "As I recall the code samples were really easy to get started with."

One survey question asked participants if they were "able to do the same research (more or less) with another data set afterwards?" – 10 of 12 respondents responded "no", indicating that for these users, there was not a viable alternative to MORF. 10 of 12 respondents also responded "high" to a question asking if "Data Utility" was "high", "medium", or "low" – the other two both responded "medium".

3.3 Challenges and Areas for Improvement

Many users also identified challenges and areas of improvement in a range of areas. First, users found it difficult to navigate the data guides (which came from the original MOOC platforms, e.g., Coursera, rather than from MORF developers) and understand where the data originally came from. One user stated, "I hate to say this, but MORF was difficult to use at the beginning, and the data guide was not very user-friendly." Another user stated that "There is no comprehensive data guide, so you need a specific expert to connect you to it for research... The variables you need are usually scattered across different tables, which can make them hard to find." A third user specifically requested a different "clear data guide" which would "explain what is in each table and how to extract data from it." A fourth user asked for "Having a codebook of all the variables in the database." Users also said that having sample

[1] Available online at https://osf.io/pjmuv/overview?view_only=b4b12a8f39084afebd20bcc50f318dd4

data values or synthetic data that they could download would have facilitated analysis.

Documentation of the MORF platform itself was another limitation identified by users. When asked "What changes to the MORF system would make it more broadly used?" two users responded that MORF designers should "Improve the documentation available within the system" and "Provide better documentation." Another user said "the main limitations I found was the not clearly documented access and usage docs for using MORF API… Due to unclear documentation, it became a challenging task".

Data extraction (pulling data from the database into a format ready for input into analysis code) was also difficult for some users. One user said "There is a steep learning curve to get familiar with it in order to create the auto-data extraction process. Once you get the hang of it, it becomes easier to use (though I think I have still explored less than 40% of the data stored in the database)". Another user recommended that MORF developers "Make data extraction easier."

The overall quality of the data provided by the MOOC vendors was also an issue. Some users reported that they observed duplication or inconsistency in tables; one user stated that "there were alot (sic) of seemly dupliciated vairables (sic) and it was always unclear to me which ones were correct and should be used" and further suggested "Having rollup tables that aggregated basic data at the student level and course levels across all courses/students. (e.g. student table: student, course, start date, end date, complete, passed, number of assignments completed…. etc.). Having example queries with variables that researchers tend to use." One user suggested "I think richer data would help. Which isn't really a criticism of MORF so much as MOOCs themselves."

Though some users liked the availability of Docker containers, another user found it "hard to work with the docker containers." Even after using the platform, two users expressed uncertainty as to what Docker containers were, suggesting that they may have used the code examples sufficiently smoothly to not have had to engage with Docker directly. These responses indicate a variety of levels of expertise among the MORF users. Unfortunately, while we asked about educational level in the survey, we did not ask about participants' levels of machine learning or programming experience and were therefore not able to explicitly compare responses in terms of this.

One user noted that better support for debugging their code was needed: "there were a lot of issues with bugs that I couldn't solve myself and getting in touch with someone who had access to the servers to debug was extremely difficult." This issue had to do with the difficulty of debugging without access to unrestricted outputs; any outputs beyond a standard set of approved libraries, including debugging code, had to be checked and approved by a member of the core development team, who was sometimes unavailable. This user also said that, contrary to other respondents, they were able to accomplish "not much" when using the platform. Another user said that "in the initial days of my usage you couldn't really see an error message! If you code failed it just did" while another said "decoding some error messages (e.g., internal server error)… would be extremely useful for researchers" and a third said "also, the error logs for MORF are hard to decode."

Relatedly, several users recommended increasing the number of output functions and tools for exploring the model. This was experienced by the development team as a continual challenge as well, to review code, consider possible misuses and risks, for making a broad number of output functions available.

Database downtime was brought up by one user, who stated, "The main obstacle for me was the connectivity issues I encountered.", though another user treated this as a limitation of a specific version: "A couple of years ago, it was almost impossible just to establish communication with the most recent database… and conduct any analysis."

4 Discussion and conclusions

Taken together, the findings from this user survey both corroborate much of the post-mortem perspective on MORF offered by Baker and Hutt [1], but also diverge in key ways. There was substantial agreement between developers and users that the platform's prioritisation of privacy protection introduced significant friction into everyday research workflows, ultimately reducing usability. Both developers and users identified difficulties related to debugging, limited visibility into data, restrictions on outputs, and how these limitations disrupted the normal iterative cycle of analysis. Users' accounts of delayed or opaque error messages, constrained diagnostic outputs, and dependence on the availability of technical support from the development team all echo developers' concerns that MORF's safeguards imposed costs that were often not adequately mitigated by either tool support or documentation [e.g. 1].

At the same time, the user perspective reveals several important differences in emphasis. In particular, the developers appear to have overestimated the degree to which Docker itself was perceived as a primary barrier – a divergence consistent with the long-recognised tendency for those who build a system to misjudge which of its aspects users will find most difficult [7]. While some users did report difficulty working with Docker containers, others described Docker as a strength, noting its flexibility, portability, and support for multiple programming languages. Notably, the users accessed MORF prior to widespread availability of Artificial Intelligence in Integrated Development Environments (e.g., Antigravity, Codex, Claude Code). This current availability may reduce barriers to using systems that require Docker Containers and multiple coding languages. Two users were able to rely on the provided examples without engaging deeply with containerisation at all, suggesting that Docker functioned as a largely invisible infrastructure layer for these researchers rather than an obstacle. This pattern is characteristic of infrastructure more broadly, which tends to recede from view until it breaks down and becomes visible as an object of concern [27]. The core usability challenges of MORF were less about containerisation per se than about the surrounding ecosystem that became salient when analyses failed to function – documentation quality, data guides, debugging affordances, and the clarity of system feedback. This finding highlights the importance of collecting user feedback systematically (as done here) rather than relying solely on informal reports, which may disproportionately represent frustrated users. However, another possible interpretation of this finding is that the users were more positive about Docker because

many researchers who were uncomfortable with the idea of using Docker did not attempt to use the platform at all. These individuals would not have shown up in a survey of past users. As such, limitations in the platform's usability not only shaped how effectively researchers could accomplish their goals, but also played a major role in determining who was able to use the platform for research at all.

Users were also more positive than the developers about the value proposition of the platform itself. A majority of respondents reported that MORF enabled research that they could not have conducted otherwise, and half produced publishable outputs using MORF data. This indicates that despite the substantial barriers to use, the uniqueness and scale of the data were sufficient to justify the effort for some researchers – particularly graduate students and early-career researchers without the resources to collect their own data sets, or seeking affordable (in this case free) access to otherwise unavailable datasets (explicitly mentioned by several participants), or investigating research questions uniquely addressable with MORF data (also mentioned by several participants). However, the fact that successful projects often required persistence and support from the development team (the human infrastructure that, in practice, often makes technical infrastructure usable [14]), as well as the very limited number of past users available to survey, reinforces the developers' viewpoint that MORF did not scale well to independent external use.

These findings suggest several design implications that developers of privacy-preserving research infrastructure must address if such platforms are to serve a broad community of researchers, in education and beyond. First, enclaves must support iterative development and debugging workflows that more closely approximate standard data science practice, even under strict privacy constraints. This includes clearer, more informative error reporting, tools that enable fast feedback cycles, and privacy-preserving mechanisms for exploratory analysis and diagnostics. As part of this, a broader range of output functions – and faster processes for validating new output functions – are essential, as both developers and users noted. Second, comprehensive, researcher-oriented documentation is also needed, both for the enclave and for the underlying data. This factor was more heavily emphasised by the users than by the developers. Users also discussed the need for clear codebooks, example queries, sample values, synthetic data examples, more streamlined data structures, cleaner and more consistent data, and higher-level "roll-up" tables to reduce the cognitive and technical burden of data discovery. While many of these can be provided by a platform, limits in documentation and data quality go beyond what developers can address by themselves and indicate a need for greater partnership between enclave developers and the learning platforms contributing data to maximise the potential benefit for researchers [20]. Established frameworks such as the FAIR principles (findable, accessible, interoperable, reusable) offer a useful starting point for structuring such partnerships and guiding shared standards for data documentation and quality [20].

Finally, enclave designers must explicitly plan for user independence at scale. Platforms should minimise reliance on informal developer intervention (for instance, to interpret error messages and check debugging output for PII) and instead embed better support, guardrails, and guidance directly into the system. Human technical support may be warranted in the early stages of platform development, but this approach creates frictions even when sufficiently resourced and is unsustainable in the longer term. For example, future enclaves could invest in automated privacy-preserving error reporting systems that can intelligently redact sensitive information from error messages while preserving diagnostic value (cf. [26]). A complementary direction targets the documentation and data-discovery burden that users identified as a central obstacle. Community mechanisms such as forums in which researchers share workflows, example queries, and hard-won knowledge of the data - could make it easier for researchers to pick up the expertise that MORF users had to obtain individually from the development team, formalising part of the human infrastructure the platform implicitly relied on [14]. This accumulated knowledge might in turn form a natural basis for creating a human-AI collaboration tool, such as an assistant that learns from shared workflows and prior queries/responses, and then helps researchers locate relevant variables, interpret opaque error messages, and navigate unfamiliar data structures [e.g. 29].

Overall, while the example of MORF shows how difficult it can be for enclaves to reduce barriers for users without compromising privacy, the success of enclaves in other fields [8, 21] suggests that careful attention to usability and researcher workflows can meaningfully shift the balance. There are many cases where direct data sharing is infeasible, and enclaves represent an alternative way for researchers to access the data they need. MORF's experience demonstrates that researchers are willing to work within privacy-protecting constraints when the data and research opportunities are sufficiently valuable. If sufficient usability can be achieved, and the value of the data is clear, a broader community of researchers can be empowered to conduct important research on sensitive data, while protecting the privacy of the learners and instructors (and other types of participants, in non-educational domains) whose data makes the research possible.

Acknowledgments

We would like to thank the developers of the MORF platform for their support for this research, the users of the MORF platform for their responses, the National Science Foundation for funding support, and Chelsea Porter for assistance in document preparation. Claude Opus 4.8 was used for document preparation and editing, but all text remains the responsibility of the authors.

References

[1] Ryan S. Baker and Stephen Hutt. 2025. MORF: A Post-mortem. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference* (LAK '25). Association for Computing Machinery, New York, NY, 797–802. https://doi.org/10.1145/3706468.3706478

[2] Neil Bosch, Robert W. Crues, Najmuddin Shaik, and Luc Paquette. 2020. "Hello, [REDACTED]": Protecting Student Privacy in Analyses of Online Discussion Forums. In *Proceedings of the 13th International Conference on Educational Data Mining* (EDM 2020), A. N. Rafferty, J. Whitehill, C. Romero, and V. Cavalli-Sforza (Eds.). International Educational Data Mining Society, 39–49.

[3] Judy Bowen and Annika Hinze. 2024. Shared Data Visualisations for Private Data: Managing Ownership and Visibility. In *Proceedings of the 36th Australasian Conference on Human-Computer Interaction* (OzCHI '24). Association for Computing Machinery, New York, NY, 213–222. https://doi.org/10.1145/3726986.3727010

[4] Virginia Braun and Victoria Clarke. 2021. *Thematic Analysis: A Practical Guide.* SAGE Publications, London.

[5] Rungson Chomeya. 2010. Quality of Psychology Test between Likert Scale 5 and 6 Points. *Journal of Social Sciences* 6, 3, 399–403. https://doi.org/10.3844/jssp.2010.

399.403
[6] Family Educational Rights and Privacy Act of 1974 (FERPA), 20 U.S.C. §1232g.
[7] John D. Gould and Clayton Lewis. 1985. Designing for Usability: Key Principles and What Designers Think. *Commun. ACM* 28, 3, 300–311. https://doi.org/10.1145/3166.3170
[8] Melissa A. Haendel, Christopher G. Chute, Tellen D. Bennett, David A. Eichmann, Justin Guinney, Warren A. Kibbe, Philip R. O. Payne, Emily R. Pfaff, Peter N. Robinson, Joel H. Saltz, Heidi Spratt, Christine Suver, John Wilbanks, Adam B. Wilcox, Andrew E. Williams, Chunlei Wu, Christian Blacketer, Robert L. Bradford, James J. Cimino, and Ken R. Gersing. 2021. The National COVID Cohort Collaborative (N3C): Rationale, Design, Infrastructure, and Deployment. *Journal of the American Medical Informatics Association* 28, 3, 427–443. https://doi.org/10.1093/jamia/ocaa196
[9] Langdon Holmes, Scott Crossley, Harshvardhan Sikka, and Wesley Morris. 2023. PIILO: An Open-Source System for Personally Identifiable Information Labeling and Obfuscation. *Information and Learning Sciences* 124, 9/10, 266–284. https://doi.org/10.1108/ILS-04-2023-0032
[10] Stephen Hutt, Ryan S. Baker, Michael Mogessie Ashenafi, Juan Miguel Andres-Bray, and Christopher Brooks. 2022. Controlled Outputs, Full Data: A Privacy-Protecting Infrastructure for MOOC Data. *British Journal of Educational Technology* 53, 4, 756–775. https://doi.org/10.1111/bjet.13231
[11] Aytaj Ismayilzada, Mirka Saarela, and Ayaz Karimov. 2025. Mining for Knowledge, Not Trouble: GDPR's Impact on Educational Data Mining. In *Proceedings of the 18th International Conference on Educational Data Mining*. International Educational Data Mining Society, 608–612. https://doi.org/10.5281/zenodo.15870232
[12] David Johnson, Mohammad Mushtaq, Nishaad Rao, and Noura Insolera. 2022. Updating the Geospatial Data in the PSID Restricted Data Enclave for 2005–2017. *Panel Study of Income Dynamics*. Retrieved from https://psidonline.isr.umich.edu/publications/Papers/tsp/2022-01_Geocode_Update.pdf
[13] Julia Lane and Claudia Schur. 2010. Balancing Access to Health Data and Privacy: A Review of the Issues and Approaches for the Future. *Health Services Research* 45, 5p2, 1456–1467. https://doi.org/10.1111/j.1475-6773.2010.01141.x
[14] Charlotte P. Lee, Paul Dourish, and Gloria Mark. 2006. The Human Infrastructure of Cyberinfrastructure. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work* (CSCW '06). Association for Computing Machinery, New York, NY, 483–492. https://doi.org/10.1145/1180875.1180950
[15] Erika McCallister, Tim Grance, and Karen Scarfone. 2010. *Guide to Protecting the Confidentiality of Personally Identifiable Information*. Diane Publishing.
[16] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72, 1–23. https://doi.org/10.1145/3359174
[17] Daniel Meza, Basil Khuder, Joseph I. Bailey, Sharon R. Rosenberg, Ravi Kalhan, and Paul A. Reyfman. 2021. Mortality from COVID-19 in Patients with COPD: A US Study in the N3C Data Enclave. *International Journal of Chronic Obstructive Pulmonary Disease* 16, 2323–2326. https://doi.org/10.2147/COPD.S318000
[18] Abhishek Nagaraj and Matteo Tranchero. 2023. How Does Data Access Shape Science? Evidence from the Impact of US Census's Research Data Centers on Economics Research. Working Paper No. w31372. National Bureau of Economic Research. https://doi.org/10.3386/w31372
[19] Jakob Nielsen. 1994. *Usability Engineering*. Morgan Kaufmann.
[20] Vasyl P. Oleksiuk, Oleg Spirin, and Kateryna Osadcha. 2026. A Standardisation Model of FAIR Data for Educational Sciences. *Educational Analytics of Ukraine* 2, 91–103.
[21] Beth A. Plale, Eleanor Dickson, Inna Kouper, Samitha Harshani Liyanage, Yu Ma, Robert H. McDonald, John A. Walsh, and Sachith Withana. 2019. Safe Open Science for Restricted Data. *Data and Information Management* 3, 1, 50–60. https://doi.org/10.2478/dim-2019-0005
[22] Richard Platt and Tracy Lieu. 2018. Data Enclaves for Sharing Information Derived from Clinical and Administrative Data. *JAMA* 320, 8, 753–754. https://doi.org/10.1001/jama.2018.9342
[23] Privacy Act 1988 (Cth) (Australia).
[24] Privacy and Other Legislation Amendment Act 2024 (Cth) (Australia).
[25] Priscilla M. Regan and Jolene Jesse. 2019. Ethical Challenges of Edtech, Big Data and Personalized Learning: Twenty-First Century Student Sorting and Tracking. *Ethics and Information Technology* 21, 3, 167–179. https://doi.org/10.1007/s10676-018-9492-2
[26] Kiavash Satvat, Maliheh Shirvanian, Mahshid Hosseini, and Nitesh Saxena. 2020. Crepe: A Privacy-Enhanced Crash Reporting System. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy* (CODASPY '20). Association for Computing Machinery, New York, NY, 295–306. https://doi.org/10.1145/3374664.3375738
[27] Susan Leigh Star and Karen Ruhleder. 1996. Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. *Information Systems Research* 7, 1, 111–134. https://doi.org/10.1287/isre.7.1.111
[28] Ruoxi Sun, Hanwen Wang, Hsiang-Ting Chen, and Minhui Xue. 2024. Privacy in Motion: Implementing Differential Privacy for User Motion in VR. In *Proceedings of the 36th Australasian Conference on Human-Computer Interaction* (OzCHI '24). Association for Computing Machinery, New York, NY. https://doi.org/10.1145/3726986.3727017
[29] Dakuo Wang, Justin D. Weisz, Michael Muller, Parikshit Ram, Werner Geyer, Casey Dugan, Yla Tausczik, Horst Samulowitz, and Alexander Gray. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 211, 1–24. https://doi.org/10.1145/3359313
[30] Elise M. Young. 2015. Educational Privacy in the Online Classroom: FERPA, MOOCs, and the Big Data Conundrum. *Harvard Journal of Law & Technology* 28, 2, 549–592.
[31] Andres F. Zambrano, Shreya Singhal, Maciej Pankiewicz, Ryan S. Baker, Chelsea Porter, and Xiner Liu. 2025. De-identifying Student Personally Identifying Information in Discussion Forum Posts with Large Language Models. *Information and Learning Sciences* 126, 5/6, 401–424. https://doi.org/10.1108/ILS-11-2024-0156
[32] Elana Zeide and Helen Nissenbaum. 2018. Learner Privacy in MOOCs and Virtual Education. *Theory and Research in Education* 16, 3, 280–307. https://doi.org/10.1177/1477878518815340