

University Policies on Large Language Models: Data Privacy, Security, and Governance

Heather Dorrell

University of Virginia

EDIS 7076: Technology, Learning Systems, and Culture

Instructor: Dr. Jennifer Maddrell

April 13, 2025

Abstract

The integration of Large Language Models (LLMs) in higher education presents both opportunities and challenges for universities, particularly regarding privacy, security, and institutional policy development. As AI-powered tools become more prevalent in academic settings, concerns about data protection, institutional policies and governance, and a lack of communication among stakeholders have become apparent. Many institutions lack clear policies on AI-driven data collection, storage, and security, leaving faculty and students vulnerable to risks such as data breaches, unauthorized access, and ethical misuse.

This literature review explores the institutional policies and best practices surrounding LLM adoption in higher education, focusing on risk management, faculty preparedness, and policy enforcement. Key themes include current data security and privacy concerns, institutional governance, and recommendations to mitigate risks. Findings highlight a need for stronger institutional oversight, AI-specific privacy policies, and enhanced faculty training programs to minimize security threats and ensure responsible AI integration. Additionally, the study identifies gaps in empirical research on university AI policies and compliance mechanisms.

The review concludes with recommendations for closing these knowledge gaps, including comprehensive policy development, risk assessment protocols, and multidisciplinary collaboration. The discussion will also reveal critical information gaps that hinder informed decision-making regarding LLM security and privacy.

Keywords: AI governance, university policy, student privacy, data security, faculty training, higher education compliance

University Policies on Large Language Models in Higher Education: Data Privacy, Security, Governance, and Policy

Rise of Large Language Models in Higher Education

The rise of artificial intelligence (AI) in education, particularly large language models (LLMs), has sparked both excitement and apprehension among educators and institutions. Since 2022, generative AI (GenAI) has advanced rapidly, compelling schools and instructors to reassess traditional learning models and adapt to emerging technological capabilities (Moore & Lookadoo, 2024). As AI continues to permeate various sectors, universities face the challenge of integrating these tools in a way that enhances education while addressing concerns related to academic integrity, student privacy, and pedagogical effectiveness.

Technological innovations have historically disrupted academic institutions, reshaping how knowledge is created, accessed, and evaluated. As Wright and Sarker highlight in Dwivedi et al. (2023), tools such as calculators, email, and search engines have significantly influenced academia in ways once considered controversial (Moore & Lookadoo, 2024). Similarly, LLMs are positioned at the center of a paradigm shift, necessitating a critical examination of their role in education. While early resistance to new technologies is common, past experiences suggest that complete prohibition is neither practical nor beneficial in the long run. The debate over Wikipedia as an educational tool, for instance, mirrors current discussions about AI as a collaborative assistant in learning environments. Many educators who once resisted Wikipedia's use now recognize its value, much like how AI tools are becoming difficult to ignore in academic settings (Moore & Lookadoo, 2024).

Furthermore, according to a 2024 report from Microsoft Corporation and LinkedIn, the workplace landscape is evolving as employees push for adoption of AI in their daily tasks. In their annual report with data compiled from a six-month randomized control trial of 60 Copilot customers across industries, the study observed that 75% of global knowledge workers are using AI at work already and 66% of leaders say they would not hire someone without AI skills.

As Moore and Lookadoo (2024) note, students and businesses are already integrating AI tools, signaling that universities must align with these technological trends rather than resist them. Likewise, Chen et al. (2025) assert that education institutions have a responsibility to keep pace with industry needs, ensuring that students are equipped with the skills to thrive in AI-enhanced work environments.

This paper examines the privacy and security risks associated with Large Language Models being used in public higher education settings and will consider those risks in relation to existing university governance and policies. Recommendations for policies and training are offered to manage those risks. The discussion will include recommendations for future research.

Data Privacy and Security

Universities manage vast amounts of sensitive information, from financial records to research data, making them prime targets for cybercriminals (Fouad, 2021). The integration of Large Language Models (LLMs) and Generative AI (GenAI) in higher education presents significant privacy and security challenges, particularly as these technologies interact with student-generated data.

Data Collection and Retention

LLMs present a unique challenge from other educational technologies because of how they are trained, how they operate, and how they interact with data in real time. They are generally pre-trained on massive data sets and are fine-tuned on smaller sets that may use input logs. These data sets could have personal information included. Universities do not fully own or control the AI models they use, relying on third-party vendors for assurances on data retention, fine-tuning practices, and compliance with privacy laws. As datasets are aggregated and repurposed for machine learning applications, questions arise about the transparency of data usage and the extent to which informed consent is properly obtained (Holmes et al., 2021). For instance, data deletion is a challenge, as the memorization capabilities of LLMs complicate the removal of outdated or sensitive student information, raising compliance concerns (Neel & Chang, 2024).

Data Usage

The increasing implementation of AI-driven technologies raises questions about accountability in data use. As AI continues to shape higher education, institutions must navigate the balance between leveraging AI-driven efficiencies and protecting student rights. The complexity and rapid advancement of AI technology makes it challenging to determine responsibility for decisions and actions, necessitating that institutions address ethical concerns, including potential discrimination, bias, and stereotypes while ensuring data security (Chan, 2023).

A growing area of concern involves how third-party AI vendors manage and secure institutional data. While AI-driven interventions may enhance student success, they also create risks of surveillance and data commodification, especially when institutions partner with third-party AI vendors (Chan, 2023).

LLMs trained on extensive datasets may contain sensitive and personal information, making them susceptible to data extraction attacks and privacy breaches (Wang et al., 2024). The aggregation of open-source datasets can also pose risks, as unidentified data may be discovered through behavioral proxies, leading to potential privacy violations (Holmes et al., 2021). Stakeholders may not always be aware that their textual data is being used in LLM-based automation, as consent processes are often embedded within broader enrollment agreements (Yan et al., 2023). Without effective security policies, universities face not only legal risks but also the erosion of student trust in AI-driven academic interventions (Tzimas & Demetriadis, 2021). Universities have demonstrated a varied response to third-party vendors, as among universities with ChatGPT policies, 67% have integrated the tool into teaching and learning, while a smaller portion have restricted its use (Moore & Lookadoo, 2024). There are unknown and varied reasons for these inconsistencies in application and use. For instance, financial constraints may further exacerbate disparities, as access to state-of-the-art AI models is often limited to well-funded institutions (Yan et al., 2023).

The question of data ownership in AI-driven education further complicates institutional responsibilities. Researchers and institutions must consider not only the origin of data but also how ownership rights extend beyond initial data collection. As AI models evolve through iterative learning, data stewardship must include long-term containment policies, ongoing access considerations for data providers, and communication strategies for upholding data rights beyond initial research projects (Holmes et al., 2021). This responsibility requires institutions to engage in continued oversight and collaboration with archival repositories to ensure ethical data use.

Shoshana Zuboff (2019) describes the exploitative nature of data surveillance, framing it within the concept of "surveillance capitalism," which threatens human autonomy and equity (Holmes et al., 2021). These concerns are particularly pressing for marginalized groups, such as disabled individuals, who often face significant privacy vulnerabilities. Existing privacy measures, such as de-identification, fail to adequately protect those with unique needs. If a student is the only one in a university requiring a sign language interpreter or a specialized assistive device, their identity can still be inferred despite anonymity safeguards (Gupta & Treviranus, 2020; Holmes et al., 2021). The reality for many disabled individuals is that they must trade privacy for access to essential services. Given the rising integration of GenAI tools in academic settings, universities are called to prioritize the development of ethical frameworks, institutional policies, and effective security measures to protect student data and privacy rights.

While LLMs offer numerous benefits, they are also susceptible to security vulnerabilities such as jailbreaking attacks, data poisoning, and personally identifiable information (PII) leakage (Das, Amini, & Wu, 2024). Jailbreaking, for instance, involves crafting unique inputs to trick a LLM to bypass constraints like revealing private or sensitive information. Data poisoning is the act of injecting false, biased, or malicious data into AI systems to degrade the model's integrity or manipulate the model into exposing private or sensitive data.

One of the primary ethical challenges in the adoption of AI is reconciling the need for data-driven decision-making with students' rights to privacy (Tzimas & Demetriadis, 2021). The expansion of AI in education has exacerbated the 'datafication' of student information, requiring careful consideration of the risks associated with placing personal data in institutional databases. Once student information is stored, it can be linked to other datasets, generating patterns and correlations that could have unintended consequences (Holmes et al., 2021). For instance, while data analytics can provide insights that improve education, it also raises concerns about how student performance and attendance records might be repurposed by external agencies.

Despite these concerns, historical events such as the COVID-19 pandemic have highlighted the societal benefits of data sharing. Personal data has been used in public health and medical research to safeguard communities, demonstrating the potential value of information when applied responsibly (Holmes et al., 2021). In education, similar trade-offs must be considered to determine how much personal information should be used to enhance student outcomes while maintaining ethical standards. AI adoption in higher education requires a structured approach that ensures student privacy protection remains central to institutional decision-making while maximizing AI's potential benefits.

Data Considerations for Higher Education Institutions

Without continuous and robust institutional oversight, concerns about privacy protection remain. For instance, many universities rely on cloud-based AI services where student data is stored for a prolonged period, which introduce additional security risks (Chen et al., 2024). Attackers know cloud platforms serve as a centralized repository for valuable student data, and there are many potential points of entry to gain access to the cloud. Misconfigurations of the settings, weak identity and access protections, and vulnerabilities in the integration of AI tools into Learning Management Systems are a few of those vulnerabilities. Privacy protection mechanisms are embedded in many AI tools, but there is no guarantee that training data will remain secure, making them susceptible to jailbreak exploits and unauthorized data retrieval (Wu et al., 2023). Additionally, the lack of standardized approaches for

detecting and mitigating ‘dark patterns,’ deceptive design techniques that manipulate users into compromising their privacy, represent a critical flaw in university security frameworks (Lachheb et al., 2023). Dark patterns in a data security and privacy context might include auto-opting students into using AI-driven interfaces, outputs that express confidence where it is not warranted, or coercing students into sharing personal information.

Balancing security with data accessibility is a key consideration in AI data management. Privacy-preserving techniques such as differential privacy and anonymization contribute to securing student data while also supporting learning analytics and institutional decision-making (Tzimas & Demetriadis, 2021). Universities establish retention and deletion policies to regulate data storage while maintaining the operational integrity of AI-driven educational tools. However, unlike other sectors, some universities retain data indefinitely, increasing their vulnerability to breaches that can affect former students, faculty, and research collaborators (Fouad, 2021).

The implementation of AI technologies in education has introduced discussions regarding surveillance and automated decision-making. The use of facial recognition and emotion-detection technologies has raised concerns about the balance between security and privacy, particularly in contexts where regulation of such tools is still developing (Holmes et al., 2021). While existing privacy-preserving mechanisms such as differential privacy and federated learning provide some security guarantees, they remain insufficient for large-scale university AI applications due to evolving threats and the scale of data processing (Neel & Chang, 2024). The extent to which these measures are implemented is not well documented across institutions, highlighting the need for further research into university-specific AI risk management strategies.

Existing AI security frameworks must evolve to address unique vulnerabilities posed by LLMs (Das et al., 2024). These risks underscore the need for higher education institutions to develop comprehensive policies and communication strategies to safeguard student data.

Student Data Laws and Ethical Frameworks

The deployment of large language models (LLMs) in higher education raises questions about the adequacy of existing university policies in addressing privacy, security, and ethical concerns. Holmes et al. (2021) observe that no universally accepted guidelines or regulations exist to address the specific ethical concerns posed by AI in education. Nor are there guidelines for monitoring AI compliance, leaving institutions vulnerable to potential violations of privacy laws such as the Family Educational Rights and

Privacy Act (FERPA) and General Data Protection Regulation, or GDPR (Das et al., 2024). Institutions that do not implement purposeful security measures in AI adoption run the risk of noncompliance with key privacy regulations, including FERPA and GDPR (de Fine Licht, 2023). To compound the complexity of this issue, adapting laws such as FERPA and GDPR to AI-driven learning environments remains difficult due to the evolving nature of AI capabilities and the security vulnerabilities they introduce.

FERPA is a U.S. federal law passed in 1974 that was recently revised in 2021 that protects the privacy of student education records. To maintain best practices in relation to FERPA, university employees must consider that students have the right to consensual use of their data. When universities use AI tools, LLMs, or cloud platforms, they must ensure that student data is not disclosed improperly or used in ways that violate FERPA. If AI models store or reproduce personally identifiable information (PII), the students did not consent to data use, or the third-party vendor is found to lack FERPA compliance, institutions can lose trust with students, lose federal funding, and be subject to litigation. Some institutions have taken measures to comply with FERPA by prohibiting students from submitting their work to external GenAI tools and plagiarism detection software (Moore & Lookadoo, 2024). FERPA, originally designed to protect student records, does not explicitly address AI-related risks, leading to varying institutional interpretations of its applicability to GenAI tools.

GDPR is another legal framework that sets guidelines for the collection and processing of personal data of individuals within the European Union (EU) and the European Economic Area (EEA). It also applies to non-EU organizations (such as universities or companies in the U.S.) if they process the personal data of EU/EEA residents. GDPR establishes strict data protection laws affecting universities handling EU student data, yet LLMs in U.S. institutions may not fully align with GDPR's consent and data minimization principles.

GDPR is stricter than FERPA in several key ways: it protects all types of personal data, not just educational records. GDPR requires active, informed consent, not just institutional control over record access. GDPR gives individuals the right to be forgotten, prohibits certain types of automated decision-making without human review, and places legal responsibility on institutions to demonstrate compliance through documentation and audits. If AI systems collect or store PII without consent, or if students' data is used for secondary purposes (like retraining models) without their knowledge, the university may be in violation of GDPR.

GDPR's consent requirements remain a concern in contexts where students may not have a meaningful opportunity to opt out of AI-based educational tools, especially when widespread institutional adoption limits alternatives (Holmes et al., 2021). Additionally, GDPR's transparency mandates require that universities disclose how AI tools process student data, though some scholars argue that this is challenging given the opaque nature of many LLM models (Tzimas & Demetriadis, 2021). To mitigate this risk, some universities have implemented strict internal data policies, required vendor contracts with GDPR-compliant clauses, or restricted AI tools from storing user input (Moore & Lookadoo, 2024).

While some student data privacy concerns are addressed by these laws, Holmes et al. (2021) observe that no universally accepted guidelines or regulations exist to address the specific ethical concerns posed by AI in education. Furthermore, discussions on AI's power imbalances in education raise questions about whether faculty and students can realistically opt out of institution-wide AI adoption, reinforcing calls for ethical guardrails beyond legal compliance (Holmes et al., 2021).

Some experts propose additional cybersecurity and ethical AI frameworks as guidance for universities. The National Institute of Standards and Technology (NIST) Cybersecurity Framework serves as a resource for institutions seeking to align AI implementations with FERPA's data security requirements. NIST's guidelines emphasize access control and threat detection, which some researchers view as critical to mitigating LLM-related data breaches.

Similarly, the Educause Higher Education Information Security Council (HEISC) Guidelines assist universities in assessing AI-driven cybersecurity risks and aligning AI deployments with FERPA and GDPR data protection principles. Although the HEISC Guidelines and the NIST Cybersecurity Framework provide institutions with cybersecurity risk management strategies, these frameworks do not address AI-specific vulnerabilities (Das et al., 2024). Some reports emphasize the role of cross-agency data governance in providing accountability for AI data management (Data Quality Campaign, 2023).

AI Ethics and Responsible Use Guidelines, such as those developed by the Institute of Electrical and Electronic Engineers and UNESCO, provide frameworks that help universities balance compliance with ethical considerations. These guidelines focus on AI transparency, bias mitigation, and student autonomy. Some experts suggest that aligning AI governance policies with these ethical frameworks assists universities in creating an AI ecosystem that adheres to both FERPA and GDPR mandates while addressing broader social considerations (Holmes et al., 2021; Tzimas & Demetriadis, 2021).

At the European level, the High-Level Expert Group on AI (HLEG on AI) developed the Ethics Guidelines for Trustworthy AI, emphasizing AI that is legal, ethical, and dependable (Holmes et al., 2021). These guidelines advocate for a values-by-design approach, where compliance with existing legal frameworks, respect for ethical principles, and technical stability are embedded in the system from inception rather than being an afterthought. The guidelines outline seven core requirements for Trustworthy AI, including human oversight, technical safety, privacy and data governance, transparency, non-discrimination, societal well-being, and accountability (Holmes et al., 2021). Though non-binding, these guidelines have the opportunity to influence how institutions structure AI policies.

Beyond federal and international regulations, state legislators in the United States introduced additional measures that expand upon FERPA. Reports indicate that since 2013, 47 states have enacted laws addressing student data privacy, suggesting that FERPA alone may not sufficiently regulate AI's role in education (Data Quality Campaign, 2023). Some policy analysts emphasize the importance of institutional guardrails to protect student information, arguing that universities must establish strong data governance structures (Data Quality Campaign, 2023). Others suggest that compliance requires continuous assessment and refinement to keep pace with evolving AI risks (Data Quality Campaign, 2023). The U.S. Department of Education discusses the importance of digital literacy, recommending that students be made aware of AI's risks, including surveillance and bias, to promote responsible AI adoption (Chen et al., 2025).

Universities are positioned as key actors in safeguarding student data in AI-driven education. Ensuring compliance with legal and ethical standards involves maintaining high-quality datasets, establishing policies for data ownership, and enforcing governance mechanisms. However, the implementation of AI governance policies across universities remains inconsistent, with varying degrees of regulation and enforcement. In a 2023 survey of college officials, only 32% had established policies addressing AI, and nearly half had not engaged in discussions with faculty or students regarding GenAI governance (Moore & Lookadoo, 2024). The absence of AI-specific privacy guidelines leaves faculty and students vulnerable to security risks, as institutional policies designed for traditional digital technologies do not fully account for the unique privacy challenges posed by AI-driven data collection and analysis (Chen et al., 2025).

Barriers and Challenges to Governance and Compliance

Beyond the barriers to national and international governance, as AI-driven learning tools continue to expand, regulatory frameworks that ensure compliance with privacy laws become increasingly necessary (Lachheb et al., 2023). Ensuring AI data security in higher education requires collaboration between IT professionals, legal experts, and academic leadership. However, research suggests that these groups often operate in silos, making it difficult to develop cohesive AI security policies (Data Quality Campaign, 2023). Despite concerns about AI's impact on higher education, cybersecurity policies in universities often remain fragmented, primarily managed by IT departments rather than institution-wide strategies (Ulven & Wangen, 2021).

The development of AI policies is also shaped by instructional preferences. Aligning AI policies with institutional values and academic goals is emphasized, requiring clear communication on the professional and educational applications of GenAI (Moore & Lookadoo, 2024). Faculty members, who are frequently left to make their own decisions about AI use in classrooms, may unknowingly expose sensitive student data to security risks due to a lack of institutional guidance (Chen et al., 2025). Research indicates that faculty favor integrating AI guidelines into course syllabi, reflecting a need for direct guidance on appropriate AI use in academic settings (Chen et al., 2025).

As AI adoption increases in higher education, faculty and staff training in AI governance, data privacy, and security remains a critical yet underdeveloped area. Most universities do not provide faculty training on AI literacy, despite the growing necessity of equipping instructors with knowledge on secure and ethical AI use (Chen et al., 2025). Cybersecurity training efforts in many higher education institutions focus on IT personnel, with minimal emphasis on educating faculty, staff, and students on data security best practices (Fouad, 2021).

The lack of structured education on AI governance creates gaps in understanding privacy risks, increasing the potential for academic misconduct and unintentional data breaches. Studies suggest that many faculty members lack awareness of how LLMs process and retain data (Das et al., 2024), or how learning analytics tools collect, process, and store student data leading to unintentional privacy violations in higher education environments (Tzimas & Demetriadis, 2021). Furthermore, most faculty members are not adequately trained to identify security threats associated with LLMs, which raises the likelihood of data breaches (Wu et al., 2023).

Governance Training Recommendations

Cybersecurity training programs that focus on LLM-specific threats can help faculty and staff recognize risks related to adversarial attacks and data extraction techniques (Das et al., 2024). Best practices identified in research include limiting the input of sensitive student data into LLM-based systems and verifying AI-generated outputs for biases (Das et al., 2024). Some frameworks for AI governance also include faculty training on bias detection, data ownership rights, and ethical AI intervention practices (Tzimas & Demetriadis, 2021). Training initiatives may include hands-on exercises on adversarial attacks, prompt injection vulnerabilities, and phishing simulations related to AI-driven threats (Wu et al., 2023). Additionally, training programs may incorporate data analysis skills to ensure that faculty members can use AI-driven learning analytics while maintaining student privacy (Data Quality Campaign).

The governance aspect of AI education requires addressing issues related to academic misconduct, data privacy, transparency, and accountability (Chan, 2023). Ensuring that faculty members understand AI governance frameworks can foster responsible AI use and maintain trust within university communities. Transparency, a fundamental aspect of ethical AI, involves clarifying decision-making processes and assumptions available to stakeholders, which enhances comprehension of AI systems and related outputs (Yan et al., 2023).

Efforts to integrate AI security training into higher education have highlighted the importance of interdisciplinary collaboration. Research suggests that effective AI governance requires cooperation among policymakers, educators, and cybersecurity experts to ensure responsible AI implementation in academia (Ding et al., 2024). Ethics regarding AI in education should involve collaboration across multiple disciplines, including educational technologists, learning designers, and AI developers, working alongside ethical review committees to understand risks and develop mitigation strategies (Holmes et al., 2021).

Despite these recommendations, few universities have established comprehensive faculty training programs for AI data security. While some institutions have introduced AI-focused faculty development initiatives, there continues to be a lack of empirical research evaluating their effectiveness. Studies on universities implementing AI data security training remain scarce, leaving an open question regarding the impact of existing programs on faculty awareness and institutional data protection practices. Further research is needed to assess the strengths and weaknesses of current professional development efforts and to identify best practices for AI security education in higher education.

Without strong governance that includes a clearly communicated and robust policy regarding data security and privacy risks when using LLMs, universities may inadvertently allow AI systems to track and analyze student behavior in ways that compromise autonomy and privacy (Holmes et al., 2021). AI models trained on student data can generate predictive insights about academic performance, learning behaviors, and personal traits, leading to potential biases and ethical dilemmas (Holmes et al., 2021). Furthermore, institutional policies frequently fail to incorporate privacy ethics, leaving students at risk of data exploitation in online learning environments (Lachheb et al., 2023).

Developing AI-specific university policies is noted as a necessary step to prevent potential privacy violations and data misuse associated with AI-driven chatbots (Wu et al., 2023). Additionally, universities operating in international contexts must ensure that their data collection and processing practices comply with varying legal frameworks to address cross-border regulatory differences (Tzimas & Demetriadis, 2021). As LLM adoption continues to expand, discussions around policy amendments remain critical to maintaining compliance with privacy laws and institutional governance standards.

Discussion and Recommendations for Institutional Policy Development

After reviewing the literature on AI-specific privacy and security policies of Large Language Models in higher education settings, it is clear that several factors are contributing to potential compromises in student data privacy and security rights and ethical concerns. Overarching laws and frameworks need to be synthesized and addressed by university administrators and leaders. AI Security and Governance frameworks need to be created, regularly reviewed and updated as technology evolves. Policies need to be addressed, understood, and communicated effectively throughout the institutional bodies. Also, research must continue to address these ongoing concerns and to hone best practices related to student data practices and policies.

AI-Specific Privacy and Security Policies

Implement AI governance frameworks that address:

- **Data Collection and Retention Limits:** Clearly define what student data AI tools can process, minimize long-term storage, and enforce deletion protocols to prevent unauthorized access (Das et al., 2024).
- **Risk-Based Cybersecurity Measures:** Strengthen AI security with encryption, differential privacy, multi-factor authentication, and real-time monitoring to detect vulnerabilities (Neel & Chang, 2024).

- **Mandatory Ongoing AI Risk Assessments and Reviews:** Universities should conduct annual AI security audits to evaluate compliance gaps and emerging threats (Data Quality Campaign, 2023). Conduct continuous security reviews to identify vulnerabilities and refine AI governance strategies (Das et al., 2024).
- **Transparency and Vendor Accountability:** Require AI vendors to disclose data handling practices, bias mitigation strategies, and security measures before institutional adoption (Tzimas & Demetriadis, 2021).
- **Opt-Out Option and Policy to Protect Student Rights:** Provide students with the choice to use AI tools within the context of the university and alternatives for those who wish to forgo use.
- **Faculty and Student AI Literacy Programs:** Provide ongoing training on AI privacy risks, security best practices, and ethical considerations to reduce unintentional data breaches (Chen et al., 2025).

Best Practices for Ethical AI Governance

Establish governance structures that emphasize fairness, security, and oversight:

- **Cross-Disciplinary AI Governance Committees:** Form dedicated oversight bodies with IT, legal, faculty, and student representation to guide AI policy implementation (Moore & Lookadoo, 2024).
- **Enhanced AI Governance Structures:** AI oversight should extend beyond IT departments to involve faculty, administrators, and compliance teams for accountability (Data Quality Campaign, 2023).
- **Institutional AI Auditing:** Conduct regular audits to assess AI tools for bias, inaccuracies, and security risks (Holmes et al., 2021).
- **Equitable AI Implementation:** Work toward AI applications do not disproportionately impact marginalized groups and maintain transparency in AI decision-making by providing equal access to institution-approved AI tools (Holmes et al., 2021).
- **Comprehensive Faculty Training:** Equip educators with AI security knowledge to prevent unintentional breaches and ensure ethical AI integration in classrooms (Wu et al., 2023).

Future Research Directions on AI Data Security in Universities

Long-Term Studies on AI's Impact on Student Data Security

While AI-driven learning tools are increasingly integrated into higher education, there is a lack of longitudinal research on their impact on student data security. Future studies should:

- Analyze long-term risks of AI-powered learning analytics and their potential to expose sensitive student information over time.

- Investigate the effectiveness of current university data retention and deletion policies in AI-driven environments.
- Examine whether AI privacy frameworks evolve in response to emerging threats, such as adversarial attacks and AI-driven surveillance (Neel & Chang, 2024).

Institutional and Cross-Sector Collaboration for AI Governance

Research highlights a lack of coordination between IT departments, legal teams, and faculty regarding AI governance (Data Quality Campaign, 2023). Future research should:

- Explore best practices for cross-disciplinary AI governance committees that integrate IT, legal, faculty, and student perspectives.
- Assess the effectiveness of AI security policies in institutions with centralized vs. decentralized governance models.
- Identify challenges universities face when enforcing AI compliance across departments and propose solutions to standardize AI data security protocols.

Partnerships Between Universities, Policymakers, and AI Vendors

With universities relying on external AI platforms, there is little research on how institutions, regulators, and AI companies collaborate on data security. Key areas for further study include:

- Investigating how AI vendors comply with FERPA, GDPR, and institutional policies in handling student data (Tzimas & Demetriadis, 2021).
- Examining data-sharing agreements between universities and AI providers to assess transparency and risk management.
- Proposing a standardized framework for university-AI vendor partnerships that prioritize privacy, ethical AI use, and data protection.

Empirical Studies on AI-Driven Privacy Risks

Many AI security concerns remain theoretical due to a lack of case studies on real-world AI data breaches and privacy violations in higher education. Future research should:

- Conduct case studies of past AI-related data security incidents in universities to identify patterns and vulnerabilities (Holmes et al., 2021).
- Assess student and faculty awareness of AI data security risks and how institutional policies influence user behavior.
- Explore how AI-driven decision-making (e.g., predictive analytics) may create unintended privacy risks, especially for marginalized student populations.

By addressing these gaps, future research can provide universities with actionable insights to enhance AI governance, improve compliance with evolving regulations, and mitigate privacy risks in higher education.

Conclusion

This research highlights the urgent need for universities to establish clear AI governance policies that address privacy, security, and ethical concerns. As institutions increasingly integrate LLMs into academic settings, many lack AI-specific policies, leaving students and faculty vulnerable to data breaches, surveillance risks, and regulatory noncompliance. Existing privacy frameworks such as FERPA and GDPR do not fully address AI-driven data processing, necessitating institutional oversight and policy updates to ensure compliance and protect academic integrity.

The Importance of Data Security in AI-Powered Learning

AI-driven learning environments process large amounts of student data, increasing exposure to security risks such as data leaks, adversarial attacks, and algorithmic bias. Strong security measures, including encryption, multi-factor authentication, and real-time monitoring, must be implemented to protect sensitive educational records. Additionally, AI literacy initiatives are crucial, as many faculty and students remain unaware of how AI systems collect, process, and retain their data. Universities must integrate structured AI training programs to close these knowledge gaps and promote responsible AI engagement.

The Role of Institutional Policies in AI Governance

Universities must take a proactive approach to mitigating AI-related privacy risks by defining policies for data security, retention limits, and responsible AI use. Without clear guidelines, institutions struggle to enforce compliance, and faculty lack the necessary training to integrate AI responsibly. Additionally, fragmented governance causes IT, legal, and academic

teams operate in silos, which weakens oversight. Establishing cross-disciplinary AI governance committees and enhancing faculty and student education on AI risks are essential to fostering responsible AI use while ensuring compliance with evolving regulations.

University Responsibility and the Need for Continuous Policy Updates

As AI technology rapidly evolves, universities must ensure continuous policy updates and risk assessments to address emerging threats. Institutional policies should not remain static but should adapt to new AI capabilities through regular audits, compliance reviews, and faculty development programs. Without ongoing updates, universities risk falling behind in data protection and ethical AI governance.

Balancing innovation and student privacy is essential. While AI offers transformative potential in education, its benefits must not come at the expense of privacy violations or ethical concerns. Institutions that fail to establish effective AI governance risk legal consequences and a loss of student trust. Moving forward, universities must prioritize AI-specific privacy frameworks that safeguard student data while ensuring AI-driven advancements remain ethical, transparent, and aligned with academic values.

References

- Abd-Alrazaq, A., AlSaad, R., Alhuwail, D., Ahmed, A., Healy, P., Latifi, S., Aziz, S., Damseh, R., Alrazak, S., & Sheikh, J. (2023). Large Language Models in Medical Education: Opportunities, Challenges, and Future Directions. *JMIR Medical Education*, 9. <https://doi.org/10.2196/48291>.
- A. Shoufan, "Exploring Students' Perceptions of ChatGPT: Thematic Analysis and Follow-Up Survey," in *IEEE Access*, vol. 11, pp. 38805-38818, 2023, doi: 10.1109/ACCESS.2023.3268224.
- Asimily. (2023). *4 cyberattacks that shocked universities and colleges*. Retrieved from <https://asimily.com/blog/4-cyberattacks-universities-and-colleges/>
- Berendt, B., Littlejohn, A., & Blakemore, M. (2020). AI in education: Learner choice and fundamental rights. *Learning, Media and Technology*, 45(3), 312-324.
- Blackmon, S. J., & Major, C. H. (2023). *Inclusion or infringement? A systematic research review of students' perspectives on student privacy in technology-enhanced, hybrid, and online courses*. [British Journal of Educational Technology], [54(6)], 1542-1565 pages. [https://doi.org/\[DOI\]](https://doi.org/[DOI])
- Chan, C.K.Y. A comprehensive AI policy education framework for university teaching and learning. *Int J Educ Technol High Educ* 20, 38 (2023). <https://doi.org/10.1186/s41239-023-00408-3>
- Chen, K.; Tallant, A.C.; Selig, I. Exploring generative AI literacy in higher education: Student adoption, interaction, evaluation, and ethical perceptions. *Inf. Learn. Sci.* 2024.
- Das, B. C., Amini, M. H., & Wu, Y. (2024). *Security and privacy challenges of large language models: A survey*. *Journal of the ACM*, 37(4), Article 111.
- Data Quality Campaign. (2023). *DATA 101: A briefing book for policymakers on education to workforce data*. <https://dataqualitycampaign.org/resource/data-101-briefing-book-policymakers/>
- de Fine Licht, K. (2023, August). Integrating large language models into higher education: Guidelines for effective implementation. In *Computer Sciences & Mathematics Forum* (Vol. 8, No. 1, p. 65). MDPI.
- Dempere, J., Modugu, K., Hesham, A., & Ramasamy, L. K. (2023). *The impact of ChatGPT on higher education*. *Frontiers in Education*, 8, 1206936. [https://doi.org/10.3389/feduc.2023.1206936​;:contentReference\[oaicite:0\]{index=0}](https://doi.org/10.3389/feduc.2023.1206936​;:contentReference[oaicite:0]{index=0}).

- Douglas, M. (2023). Large Language Models. *Communications of the ACM*, 66, 7 - 7. <https://doi.org/10.1145/3606337>
- Educause. (2023). *Higher Education Information Security Council (HEISC) Guidelines: Cybersecurity and Data Protection in Higher Education*. Retrieved from <https://www.educause.edu/heisc>
- Fouad, N. S. (2021). Securing higher education against cyberthreats: from an institutional risk to a national policy challenge. *Journal of Cyber Policy*, 6(2), 137-154.
- Future of Privacy Forum. (2021). *Student privacy primer*. Student Privacy Compass.
- Gupta, M., Akiri, C., Aryal, K., Parker, E., & Praharaj, L. (2023). From ChatGPT to threatgpt: Impact of generative AI in cybersecurity and privacy. IEEE Access.
- Holmes, W., & Porayska-Pomsta, K. (Eds.). (2023). *The ethics of artificial intelligence in education: Practices, challenges, and debates*. Routledge. <https://doi-org.proxy1.library.virginia.edu/10.4324/9780429329067>
- IEEE. (2023). *IEEE Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*. Retrieved from <https://ethicsinaction.ieee.org>
- Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2024). Generative AI in Higher Education: A Global Perspective of Institutional adoption Policies and Guidelines. *Computers and Education Artificial Intelligence*, 100348.
- Lachheb, A., et al. (2023). *The role of design ethics in maintaining students' privacy*. [British Journal of Educational Technology], [54 (6)], pages 1653-1670. [https://doi.org/\[DOI\]](https://doi.org/[DOI])
- Microsoft Corporation, & LinkedIn. (2024). *AI at work is here. Now comes the hard part*. Microsoft. <https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part>
- Moore, S., & Lookadoo, K. (2024). Communicating Clear Guidance: Advice for Generative AI Policy Development in Higher Education. *Business and Professional Communication Quarterly*, 87(4), 610-629. <https://doi.org/10.1177/23294906241254786>
- National Institute of Standards and Technology. (2014, updated continuously). *Framework for Improving Critical Infrastructure Cybersecurity* (Version 1.1). Retrieved from <https://www.nist.gov/cyberframework>
- Neel, S., & Chang, P. (2023). Privacy issues in large language models: A survey. *arXiv preprint arXiv:2312.06717*.

- Ruhländer, L., Popp, E., Styliadou, M., Khan, S., & Svetinovic, D. (2024). *On the security and privacy implications of large language models: In-depth threat analysis*. 2024 IEEE International Conferences on Internet of Things (iThings), IEEE Green Computing & Communications (GreenCom), IEEE Cyber, Physical & Social Computing (CPSCom), and IEEE Smart Data (SmartData). <https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics62450.2024.00102>
- Tzimas, D., & Demetriadis, S. (2021). *Ethical issues in learning analytics: A review of the field*. [Educational Technology Research and Development], 69, pages 1101-1133.
- Ulven, J. B., & Wangen, G. (2021). A systematic review of cybersecurity risks in higher education. *Future Internet*, 13(2), 39.
- UNESCO. (2023). *Recommendation on the Ethics of Artificial Intelligence*. Retrieved from <https://www.unesco.org/en/artificial-intelligence/recommendation>
- U.S. Department of Education. (2021). *Family Educational Rights and Privacy Act (FERPA)*. Retrieved from <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>
- Wu, X., Duan, R., & Ni, J. (2023). *Unveiling security, privacy, and ethical concerns of ChatGPT*. Department of Electrical & Computer Engineering, Queen's University. <https://doi.org/10.48550/arXiv.2307.14192>
- Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gašević, D. (2023). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112. <https://doi.org/10.1111/bjet.13370>
- Zhang, L., & Xie, Y. (2023). *Data security crisis in universities: Identification of key factors affecting data breach incidents*. *Humanities and Social Sciences Communications*, 10(1), 321. <https://doi.org/10.1057/s41599-023-01757-0>
- Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.