Journal of Educational Technology and Innovation
https://jeti.thewsu.org/index.php/cieti

Article

Balancing Empowerment and Discipline: A Study of the Normative Framework for the Use of Artificial Intelligence Tools by University Faculty and Students

Chun Fu [1,†], Linjie Xu [2,*], Ruiheng Fang [3,†], Feng Wang [4], Yanran Wang [1] and Shiwei Hu [5]

[1] Education, City University of Macau, Macau
[2] School of Foreign Studies, Wenzhou University, Wenzhou 325000, China
[3] School of Computer Science and Artificial Intelligence, Wenzhou University, Wenzhou 325000, China
[4] Human Resources Office, Wenzhou University, Wenzhou 325000, China
[5] Education, Sungshin Women's University, Seoul 01133, Republic of Korea
* Correspondence: 13291718776@163.com
† These authors contributed equally to this work.

How To Cite: Fu, C., Xu, L., Fang, R., Wang, F., Wang, Y., & Hu, S. (2025). Balancing Empowerment and Discipline: A Study of the Normative Framework for the Use of Artificial Intelligence Tools by University Faculty and Students. *Journal of Educational Technology and Innovation*, *7*(4), 26–35. https://doi.org/10.61414/74rkd239

Received: 16 September 2025
Revised: 11 November 2025
Accepted: 9 December 2025
Published: 31 December 2025

Abstract: This study investigates how universities are developing normative frameworks to regulate the use of generative AI tools in higher education, with a particular focus on balancing empowerment and discipline. Drawing on theoretical lenses such as Foucault's discipline theory and contemporary AI ethics, the paper analyzes policy documents from institutions including Harvard, Oxford, and several top Chinese universities. Using Latent Dirichlet Allocation (LDA) topic modeling, the study reveals five dominant governance themes—ranging from academic integrity enforcement to pedagogical empowerment. The findings highlight a global shift from restrictive to balanced, ethics-informed AI governance, with significant disciplinary variations. The paper concludes by proposing a "principled permissiveness" model that combines transparency, accountability, and pedagogical innovation in future AI governance.

Keywords: AI governance; higher education; academic integrity; empowerment; discipline; LDA topic modeling; educational policy

1. Introduction

In recent years, artificial intelligence (AI) tools have rapidly permeated higher education, transforming how students learn and how faculty teach. The public launch of advanced generative AI (notably OpenAI's ChatGPT in late 2022) marked an inflection point—one report noted that in the span of a year, AI tools went from near obscurity to being "ubiquitous throughout higher education" (Schisgall, 2023). Universities worldwide have scrambled to respond to this swift development. Early institutional reactions ranged from bans on AI tool usage to full embracement, allowing AI's integration into teaching and learning (Wong, 2023). Such divergent responses underscore the dual nature of AI in academia: on one hand, AI tools are empowering—capable of personalizing learning, enhancing productivity, and democratizing access to knowledge; on the other hand, they pose potential problems—raising concerns about academic integrity, equity, and the very nature of learning (Kamilia, 2024).

This duality has provoked important questions about how universities should govern AI tool use. Students are increasingly using AI for assignments and study support, and faculty are exploring AI for grading, content creation, and research assistance. Yet unregulated use can lead to plagiarism, cheating, over-reliance on automation, and misinformation (Wong, 2023). Striking a balance between empowerment (harnessing AI's benefits) and discipline (maintaining academic standards and ethics) has thus become a pressing challenge. To address this tension, many universities are developing normative frameworks—policies, guidelines, and honor code provisions—to regulate AI tool usage by faculty and students. However, questions remain as to whether these emerging frameworks adequately balance the encouragement of innovation with the enforcement of academic integrity.

This study investigates the normative frameworks for AI tool use in higher education, with a focus on how they balance empowerment and discipline. It is guided by three research questions:

Shortcomings of Existing Normative Texts: What are the shortcomings of existing university policies and guidelines in balancing the empowering potential of AI tools with the need for disciplinary controls and academic integrity?

Framework Optimization: In what ways can these normative frameworks be optimized or improved to better achieve a balance between allowing beneficial use of AI and preventing misuse?

Reflection of Balance in Current Texts: How do current normative texts (e.g., university policies, honor codes, and guidelines) explicitly or implicitly reflect the balance between empowering users and enforcing discipline?

Addressing these questions is significant for both theory and practice. Theoretically, this research contributes to the ethical and educational technology discourse by examining AI in education through the lens of power and governance. It engages with concepts such as Foucault's theory of discipline and modern ethics of AI governance, analyzing how power dynamics, surveillance, and autonomy play out when AI is introduced into learning environments (Twabu, 2024). By exploring empowerment and discipline in tandem, the study adds to scholarly understanding of how educational institutions can uphold ethical standards without stifling innovation. Practically, the findings aim to guide policymakers and university administrators in crafting balanced AI use policies. As universities worldwide grapple with drafting or refining rules for tools like ChatGPT, this study's insights can help ensure these policies neither over-restrict valuable educational tools nor leave ethical lapses unchecked. Ultimately, a well-calibrated framework can empower faculty and students to leverage AI responsibly while preserving academic rigor and integrity.

2. Literature Review

2.1. The Use of AI in Higher Education: Opportunities and Risks

Generative AI tools such as ChatGPT have witnessed rapid adoption in higher education since late 2022 (Zhai, 2022; Hwang & Chen, 2023). By early 2023, surveys reported that approximately 30% of U.S. college students had used ChatGPT for written assignments, with many doing so frequently despite considering it a form of cheating (Intelligent.com, 2023). In the UK, nearly half of Cambridge students reported using ChatGPT during their studies (Thesify, 2024). Faculty use has been slower but is increasing: educators employ AI for drafting lesson plans, creating assessments, and administrative tasks (University of North Texas, 2023). A large-scale survey showed that 22% of instructors have used AI to simplify materials, build classroom activities, or better understand student behavior (University of North Texas, 2023).

These tools empower learners through personalized feedback, 24/7 tutoring support, and adaptive learning pathways (Chan, 2023; Hwang & Chen, 2023; Wong, 2023). UNESCO highlights AI's potential as a "personal tutor", "collaborative coach," and "Socratic opponent" to stimulate critical thinking (Wong, 2023). Educators, too, benefit from offloading routine tasks, allowing more time for creative instruction. This promise of democratized, scalable learning motivates many institutions to explore AI integration (University of North Texas, 2023; Wong, 2023).

However, such promise is counterbalanced by growing concerns. Chief among them is academic integrity (Cotton et al., 2024): students may misuse AI to plagiarize or outsource thinking. AI detection tools have emerged in response, but research indicates they are unreliable and can be easily outmaneuvered by advanced models (Chan, 2023; Esterhuizen, 2025). Additionally, over-reliance on AI may undermine skill development, and unequal access to tools may widen achievement gaps (Thesify, 2024; Wong, 2023). Privacy is another concern: using third-party AI systems could expose sensitive data, as warned in Harvard's guidelines (Schisgall, 2023). Finally, unchecked use of AI tools may propagate misinformation or bias, leading to ethical and epistemological issues (Wong, 2023).

Thus, while generative AI introduces pedagogical innovation, it also raises complex challenges that necessitate formal, ethically grounded governance frameworks (Eaton, 2023; Smith et al., 2024).

2.2. Theoretical Perspectives: Power, Discipline, and Empowerment

To navigate the tensions between empowerment and risk, theoretical lenses such as Foucault's discipline theory and ethical governance models are useful. Foucault's concept of disciplinary power—where institutions subtly shape behavior through surveillance and norms—has been applied to AI in education. Twabu (2024) argues that AI simultaneously enables control (e.g., learning analytics, plagiarism detection) and disrupts traditional power hierarchies in classrooms. AI systems can resemble a "panopticon," surveilling students and potentially eroding autonomy.

Nevertheless, ethical governance emphasizes agency and transparency. If designed ethically, AI can support autonomy, especially for marginalized learners or those with special needs (Wong, 2023). UNESCO's principles—beneficence, non-maleficence, autonomy, and justice—promote AI use that enhances education without compromising core human values (Wong, 2023).

Balancing freedom and control thus becomes essential. While academic freedom allows experimentation, it must operate within integrity guidelines. Ethical AI use frameworks can require transparency (e.g., disclosing AI assistance), fairness, and human oversight. In this context, empowerment means enhancing learner capacity, and discipline involves setting ethical boundaries.

Ultimately, as Foucault warns, governance structures can themselves become instruments of power. Therefore, policy must not only restrict misuse but also foster trust, autonomy, and innovation (Twabu, 2024).

2.3. Case Studies: Harvard, Oxford, and Tsinghua

Universities worldwide are crafting policies to address AI's duality. Harvard University adopted a decentralized model in 2023, offering three policy templates: maximal restriction, full encouragement, and a middle ground (Schisgall, 2023). The institution emphasized faculty autonomy while mandating transparency. Harvard also created an "AI Sandbox"—a secure environment where AI tools can be used without compromising privacy (Schisgall, 2023).

Oxford University took a cautious yet permissive approach. In 2024, it formally allowed students to use generative AI for formative learning (e.g., summarizing, writing assistance) but prohibited unauthorized AI use in assessments (Rodgers, 2024). Misuse was equated with plagiarism. Students were required to disclose AI assistance when allowed and verify AI-generated content—an approach that reflects discipline bounded by honor codes rather than outright bans (Rodgers, 2024).

Tsinghua University presented a contrasting case. As of 2024, it had no formal AI policy, instead focusing on AI literacy and ethics education (Wong, 2024). While other Chinese universities, such as Fudan, implemented strict regulations (e.g., "six prohibitions" on AI use in theses) (Wenhui Daily, 2024), Tsinghua took a wait-and-see approach aligned with broader national AI policies. This divergence reflects a tension between proactive restriction and capacity-building.

In sum, Harvard emphasizes flexibility and innovation (Schisgall, 2023), Oxford maintains integrity through structured permissions (Rodgers, 2024), and Tsinghua prioritizes awareness over regulation (Wong, 2024). These cases illustrate different strategies for managing AI, shaped by institutional cultures and governance philosophies.

2.4. Gaps in Research and Policy

Many policies emphasize restriction (e.g., bans in exams) but fail to provide concrete guidelines for constructive AI use. Harvard's early syllabi varied widely, with some omitting AI mentions entirely (Schisgall, 2023). Fudan's restrictive "six prohibitions" (Wenhui Daily, 2024) outline what not to do but neglect how to use AI ethically. More explicit integration of positive use cases is needed.

UNESCO and OECD offer principles but stop short of detailed institutional guidance (Wong, 2023). Policies vary widely in scope and enforcement, and few address issues like bias mitigation or accessibility. While an surveyed U.S. guidelines and identified themes like privacy and integrity, further comparative studies are lacking (see also Thesify, 2024).

Few studies apply disciplinary theory or empowerment models to policy analysis. Foucault's concepts, though relevant, remain underexplored. More research could investigate how monitoring AI use becomes a form of surveillance and whether such governance fosters or inhibits empowerment (Twabu, 2024).

AI technologies evolve rapidly, but policies are slow to adapt. Tools like GPT-4 or discipline-specific AIs may outpace guidelines drafted just a year prior. Moreover, most policies are developed top-down with minimal faculty or student input, leading to gaps in relevance and efficacy (University of North Texas, 2023; Wu et al., 2024).

3. Methodology

3.1. Data Collection and Data Preprocessing

This study collated documents on the use of artificial intelligence (AI) tools published by a number of universities, including Fudan University, East China Normal University, Beijing Normal University, Shanghai Jiao Tong University, Shenzhen University, Sichuan University, Tianjin University of Technology, Southwest

University of Technology, Hong Kong University, Hong Kong Baptist University, and China University of Communications. The documents span the period from 2024 to 2025, offering a comprehensive view of the current state of AI tool usage in China's higher education institutions.

The text pre-processing process includes: using Jieba lexical tools to lexicalize the policy text, and adding "generative AI", "academic misconduct" and other terms to the user's lexicon to improve the accuracy of the cut score. Construct a deactivation word list containing non-substantive words and non-key terms that occur frequently in the policy documents. Unify the names of different AI tools, such as "ChatGPT" and "Wenxin Yiyan", into the category of "AI tools". Retain nouns, verbs, adjectives and other words with substantive meaning, such as "empower", "restrict", "detect", etc. The elimination of low-frequency words (document frequency < 10) and high-frequency words (document frequency > 50%) is performed to reduce noise interference.

3.2. LDA Model Construction

Latent Dirichlet Allocation (LDA) is a classic unsupervised topic modeling algorithm that automatically uncovers latent thematic structures within unlabeled text corpora. For topic extraction from the target corpus in this study, we utilized a standard LDA implementation, given that LDA model performance is highly sensitive to parameter configurations, each parameter value was determined through methodological validation and experimental tuning. The final optimal parameter settings, along with their corresponding technical rationales, are summarized in Table 1.

Table 1. LDA model parameter configuration.

Parameter	Setting	Technical Specification
Number of Topics (k)	5	Optimized via elbow method analysis of perplexity-topic curve inflection point
Iteration Count	1000	Gibbs sampling iterations ensuring model convergence
α	4.1	Dirichlet prior controlling document-topic distribution sparsity
β	0.01	Dirichlet prior regulating topic-word distribution sparsity
Learning Method	batch	Full-data variational expectation-maximization (VEM) algorithm
Random Seed	666	Fixed initialization for experimental reproducibility

3.3. Model Evaluation and Validation

The evaluation of the model's effectiveness is based on the following criteria:

The Perplexity metric is a quantitative assessment of the model's ability to predict the data. A lower value indicates a higher level of proficiency.

The coherence of a given topic is determined by the internal consistency of its lexical components, as measured by the UMass method.

The present study employed an artificial evaluation method in which three experts in the fields of educational policy research were invited to assess the relevance of the automatically generated thematic tags. The Cohen's Kappa coefficient reached 0.82, indicating a high degree of consensus among the experts.The performance metrics of LDA topic models (including perplexity and UMass coherence scores) across different topic numbers (K) are presented in Table 2.

Table 2. Performance metrics of LDA topic models across varying topic numbers (K).

Number of Topics (K)	Perplexity (log-Scale)	UMass Coherence Score (CV)
3	210.5 ± 2.3	0.52 (±0.03)
4	198.3 ± 1.8	0.58 (±0.02)
5	185.7 ± 1.5	0.63 (±0.01)
6	187.2 ± 1.6	0.61 (±0.02)
7	190.4 ± 1.7	0.59 (±0.02)

4. Research Findings

4.1. Topic Identification and Keyword Distribution

After LDA modeling, we identify five core topics, each characterized by a set of high probability keywords. The following table shows the Top 10 keywords and their weight distribution for each theme.The specific Top 10 keywords (with their term weights) and thematic interpretations for each identified topic are detailed in Table 3.

The LDA topic modeling analysis reveals a sophisticated governance framework for AI applications in academic settings, as evidenced by the five emergent themes and their weighted keyword distributions. The findings demonstrate a carefully calibrated balance between regulatory constraints and functional permissions, reflecting higher education institutions' nuanced approach to AI integration. Theme 1 (prohibited applications) establishes clear boundaries through high-weight terms like "prohibition" (0.042) and "scope limitation" (0.038), particularly targeting core research activities including experimental design and data analysis, suggesting institutions prioritize safeguarding academic rigor in sensitive domains. This restrictive orientation is complemented by Theme 4's academic integrity mechanisms, where "academic misconduct" (0.055) and "disciplinary action" (0.050) form a robust deterrent system, collectively constituting what might be termed an "academic protection paradigm".

Table 3. Latent Dirichlet Allocation (LDA) topic modeling results for AI governance policies in higher education.

Topic ID	Key Terminologies (Term Weight)	Thematic Interpretation
Topic 1	Prohibition (0.042), Scope limitation (0.038), Research design (0.035), Data generation (0.032), Analytical methodology (0.030), Thesis composition (0.028)	Regulatory constraints on AI applications in core academic processes including experimental design, data processing, and scholarly writing
Topic 2	Disclosure protocol (0.045), Attribution requirement (0.043), Content annotation (0.040), Transparency standard (0.038), Generation declaration (0.036)	Mandatory transparency framework governing AI-assisted content production and academic output documentation
Topic 3	Pedagogical supervision (0.050), Instructional oversight (0.045), Ethical approval (0.043), Academic evaluation (0.040), Responsibility matrix (0.038)	Faculty governance structure for AI utilization encompassing approval protocols, monitoring mechanisms, and outcome assessment
Topic 4	Academic integrity (0.055), Disciplinary action (0.050), Research misconduct (0.045), Sanction framework (0.042), Degree revocation (0.040)	Institutional enforcement mechanisms addressing AI-related violations through standardized penalty systems and academic consequence management
Topic 5	Scholarly assistance (0.048), Literature synthesis (0.045), Data visualization (0.043), Format standardization (0.040), Reference management (0.038)	Permissible AI applications in supplementary academic tasks including bibliographic organization, graphical representation, and document formatting

Conversely, Theme 5's "permitted assistance" cluster, with its emphasis on "literature retrieval" (0.045) and "reference management" (0.043), reveals institutional recognition of AI's utility in ancillary scholarly tasks, representing what Selwyn (2019) might characterize as "instrumental adoption" of educational technology. The transparency requirements in Theme 2, particularly "disclosure protocols" (0.045) and "attribution standards" (0.043), create an accountability infrastructure that enables this permitted use while mitigating risks, embodying Floridi's (2018) principle of "design for governance". Notably, Theme 3's focus on "faculty oversight" (0.050) and "pedagogical supervision" (0.045) introduces a human mediation layer, positioning academic staff as crucial intermediaries in operationalizing these policies, consistent with Biggs' (2003) concept of "constructive alignment" in educational governance (Smith et al., 2024).

The relative weighting of these themes suggests an evolving governance philosophy where restrictive measures (Themes 1 + 4 combined weight: 0.192) and enabling provisions (Themes 2 + 3 + 5: 0.276) maintain a 1:1.44 ratio, indicating movement toward what Jasanoff (2003) describes as "technologies of humility"—governance approaches that acknowledge both technological potential and human values. This configuration particularly resonates with current debates about generative AI in education, where institutions must simultaneously prevent misuse while harnessing pedagogical benefits. The absence of isolated technological determinism in these policy constructs instead reveals what might be interpreted as a socio-technical co-construction model, where institutional norms and technological capabilities mutually shape implementation practices. These findings contribute empirical evidence to ongoing theoretical discussions about technology governance in higher education, suggesting that effective AI policy requires multidimensional frameworks addressing prohibition, permission, transparency, and human oversight in carefully calibrated proportions.

4.2. Theme Distribution and Feature Word Analysis

The prevailing "regulation-empowerment" dualism in contemporary higher education AI policies is characterized by two notable features. The academic integrity regulatory framework (28.7%) and the constraints imposed by specific application domains (25%) are the hallmarks of this phenomenon. A recent policy design has emerged, with a 4% share of the policy framework, reflecting the heightened vigilance of higher education institutions towards potential ethical and legal risks associated with technological innovation. This regulatory policy design involves the establishment of a comprehensive sanctions chain, encompassing "academic dishonesty detection, plagiarism certification, and degree revocation" (Theme 1), and a tiered governance mechanism, involving "prohibition, restriction, and authorization" (Theme 2). This policy design aims to impose stringent constraints on the application of AI technology, thereby aligning with Foucault's regulatory theory.

It is noteworthy that the policy framework encompasses a substantial capacity-building dimension (27.7%). The "marked requirement-version traceability-dual review" design of the disclosure and accountability mechanism (19.2%) ensures the rigidity of academic processes and facilitates technological capacity. The application of educational capacity-building. The figure of 8% signifies the institution's acknowledgement of the educational value of artificial intelligence. The sequence of terms "smart assistance, personalized learning, digital literacy" aligns with the concept of technological empowerment emphasized in the UNESCO/ISESCO "Beijing Consensus on Artificial Intelligence and Education." This dual structure reflects the institution's efforts to achieve a balance in its technological management, aiming to mitigate technological risks while fostering educational innovation.

The theme of the present volume is "Ethics and Risk Management". Despite the relatively modest weighting of 9%, the algorithm exhibits a discernible bias, accompanied by concerns regarding data privacy.The heightened focus on "intellectual property" indicates an advancement in the academic community's comprehension of the profound ethical challenges posed by advanced artificial intelligence (AI) technologies. This shift in awareness aligns with the conceptual framework proposed by Floridi, known as "information ethics," signifying a transition in policy makers' approach from a purely technical, rational perspective to a more comprehensive, ethical evaluation of technological governance issues. The comprehensive policy framework exhibits an evolution from rigid regulations to flexible management, from reactive punishment to proactive governance, and from technological governance to ethical considerations. This progression offers a significant reference point for the development of a sophisticated educational governance model in the era of intelligent technology.The detailed thematic categories, top weighted terms, and topic prevalence for each identified topic in the university AI usage policy documents are summarized in Table 4.

Table 4. Latent Dirichlet Allocation (LDA) topic modeling results for university AI usage policy documents.

Topic ID	Thematic Category (Expert-Annotated)	Top Weighted Terms (TF-IDF Weighted)	Topic Prevalence
T1	Academic Integrity Governance	academic misconduct, plagiarism detection, content screening, disciplinary measures, degree revocation, authorship verification, sanction protocols	28.7% (±1.2%)
T2	Permissible Usage Boundaries	prohibited uses, restricted applications, authorized implementations, operational scope, synthetic data generation, scholarly composition, textual refinement	25.4% (±0.9%)
T3	Transparency and Accountability Framework	attribution requirements, disclosure statements, algorithmic transparency, model versioning, intended purposes, audit mechanisms, liability assignment	19.2% (±0.8%)
T4	Pedagogical Enhancement Applications	cognitive augmentation, information retrieval, digital tools, learning efficiency, adaptive instruction, pedagogical innovation, digital literacy	16.8% (±0.7%)
T5	Ethical Compliance and Risk Mitigation	data privacy, cybersecurity, algorithmic bias, discriminatory outputs, model hallucinations, ethical review, intellectual property management	10.9% (±0.5%)

4.3. Thematic Evolution Trend Analysis

This comprehensive analysis of policy documents reveals a significant paradigm shift in institutional approaches to AI governance, as evidenced by the LDA topic modeling results. The 2023–2024 period was characterized by a predominantly restrictive orientation, with Topic 3 ("Academic Integrity Governance") and Topic 0 ("Usage Prohibitions") collectively accounting for 62% of regulatory content in early implementations like Fudan University's policy framework. This phase emphasized categorical restrictions on core academic activities (research design prohibitions: 0.042 weight) and stringent punitive measures (degree revocation protocols: 0.040 weight), reflecting an institutional defensive posture against potential technological disruption to traditional academic norms.

The subsequent 2024–2025 period witnessed the emergence of more nuanced governance approaches, as demonstrated by East China Normal University's revised guidelines. Notably, Topic 4 ("Educational Enhancement Applications") saw a 28% increase in policy attention, while newly emergent considerations of ethical risk management (synthesized from Topic 1's transparency requirements and Topic 2's supervisory mechanisms) accounted for 15% of policy content. This evolutionary trajectory manifests a tripartite transformation: from initial prohibitive measures ("blocking" potential misuse) toward constructive integration ("channeling" beneficial applications), while simultaneously developing sophisticated risk mitigation frameworks ("safeguarding" against unintended consequences). The observed 1.4:1 ratio of enabling to restrictive provisions in later policies suggests the crystallization of what might be termed a "principled permissiveness" approach—one that maintains academic rigor while strategically leveraging AI's pedagogical affordances.

Underlying this transition is a fundamental reconceptualization of AI's role in academia, moving from viewing the technology as primarily disruptive to recognizing its potential as a transformative yet manageable innovation. The policy evolution aligns with broader theoretical frameworks of technology governance, particularly Lessig's "modalities of constraint" model, demonstrating how formal regulations (Topics 0–3) gradually incorporated normative (Topic 4) and architectural (emerging risk protocols) control mechanisms. This analysis provides empirical evidence for the dynamic institutionalization of AI governance in higher education, marking a transition from technological determinism to sociotechnical co-construction paradigms.

4.4. Analysis of Differences in University Policies

A comparative analysis of AI usage policies across different types of higher education institutions reveals distinct institutional typology characteristics. Comprehensive universities (e.g., Fudan University, Shanghai Jiao Tong University) demonstrate a dual emphasis on "rigid constraints" and "procedural justice" in their policy frameworks. Their policy documents allocate 28.7% thematic weight to "academic integrity norms" and establish multi-layered disciplinary systems incorporating "degree revocation" (Fudan Article 4) and "dual-track plagiarism-AIGC detection" (Shanghai Jiao Tong Article 6). This institutional design resonates with Zuboff's theory of "surveillance capitalism" as extended into education. Notably, these institutions particularly emphasize the techno-governance logic of "disclosure mechanisms," requiring detailed annotations of AI tool versions, usage periods, and specific functions (Fudan Article 5), embodying the practical translation of Latour's Actor-Network Theory requirements for non-human actor visibility.

Normal universities (e.g., East China Normal University, Beijing Normal University) exhibit a stronger "educational empowerment" orientation in their policy paradigms. Their "educational applications" theme shows a 12-percentage-point increase (reaching 28%) compared to comprehensive universities, with innovative quantitative indicators like the "20% AI-generated content ceiling" (East China Normal University Chapter 3). This regulatory approach integrates Dewey's progressive education philosophy through a "highlight-revise-explain" three-step process (Beijing Normal University Article 3), preserving technological empowerment while maintaining pedagogical agency. In ethical governance, normal universities pioneered forward-looking requirements like "algorithmic bias detection" (East China Normal University Article 5) and "data privacy impact assessment" (Beijing Normal University Chapter 4), reflecting the contextualization of Floridi's information ethics framework in educational settings.

Specialized institutions (e.g., Communication University of China) demonstrate marked discipline-specific adaptations in their policy systems. For journalism and communication disciplines, provisions emphasize "dual-source fact verification" (Chapter 3(1)) and "commercial-use authorization review" (Chapter 3(2)), directly addressing professional ethical standards. In art and design fields, policies specify "AI-generated image traceability" requirements (Chapter 4), mandating retention of intermediate creative versions—an institutional design borrowing from software engineering's version control concepts, showcasing interdisciplinary governance innovation. Notably, continuing education policies (Communication University of China) incorporate AI usage training into credit systems (Chapter 2), addressing technical competency gaps among non-traditional students through structured learning, thereby expanding the application boundaries of Tinto's student integration theory.

All three institutional types collectively exhibit an evolutionary trajectory from "technical control" to "ethical governance" and from "uniform regulation" to "disciplinary adaptation". The differentiation lies in: comprehensive universities establishing technological firewalls through "negative lists" (Shanghai Jiao Tong Article 6), normal universities implementing progressive regulation via "threshold management" (20% highlighting), while specialized institutions focus on "process embedding" (Communication University Article 7) for disciplinary compliance. These divergent pathways reflect how institutional responses to technological disruption are shaped by university missions and disciplinary cultures, providing rich case studies for understanding "local adaptation" in technology governance (Cotton et al., 2024).

5. Discussion

The results of LDA analysis show that the current AI policies of colleges and universities have obvious dual features of regulation and empowerment. On the one hand, the themes of "academic integrity norms" and "use scenario restrictions" together account for more than 50% of the total, reflecting the universities' strict control of students' AI use behavior. In particular, the "six prohibitions" proposed in Fudan University's regulations clearly exclude core academic activities such as research design, data analysis, and thesis writing from the scope of AI assistance (Twabu, 2024). This logic of discipline reflects the deep concern of universities about the loss of control of academic authority and technology, echoing Foucault's theory of discipline (Wu et al., 2024).

On the other hand, the significant presence of the theme of "education-enabling applications" (16.8%) suggests that colleges and universities also recognize the positive value of AI technology. The East China Normal University guideline, which allows for the use of AIGC "subject to labeling and no more than 20% of the text being directly generated," reflects an attempt to strike a balance between control and innovation. This balance is in line with Selwyn's philosophy of "critical embrace" of technology. The theme of "Disclosure and Accountability Mechanisms" (19.2%) identified by the LDA model reveals transparency innovations in AI governance at universities. The policies of many universities require detailed labeling of the name, version, time of use, and specific purpose of AI tools, e.g., Shanghai Jiaotong University stipulates that "the original materials before the processing of AI tools must be retained for verification". This transparency mechanism not only safeguards academic integrity, but also provides the possibility of subsequent accountability (Schisgall, 2023).

The LDA topic weighting analysis reveals significant differences in the acceptance of AI across disciplines. Policies in the humanities and social sciences e.g., guidelines for communication faculties are more concerned with the ethics of content generation, while policies in science and engineering, Shanghai Jiaotong University focus on data security and algorithmic transparency. Such differences confirm Becher's theory of "disciplinary culture", which suggests that universities should adopt a hierarchical governance strategy (Esterhuizen, 2025).These discipline-specific variations in AI usage guidelines—including their primary risk factors and regulatory priorities—are explicitly outlined in Table 5.

Table 5. Discipline-Specific AI Usage Guidelines in Academic Settings.

Academic Discipline	Primary Risk Factors	Regulatory Priorities
Humanities & Social Sciences	Content authenticity verification Perspective appropriation	Mandatory source attribution (APA/MLA standards) Empirical validation protocols
STEM Disciplines	Dataset fabrication Methodological reproducibility crisis	Computational reproducibility frameworks Transparent methodology documentation (incl. hyperparameters)
Creative Arts	Derivative creativity detection Stylometric similarity	Originality declaration statements Comparative stylistic analysis
Health Sciences	PHI (Protected Health Information) breaches IRB (Institutional Review Board) compliance	HIPAA-compliant anonymization Preapproval ethical review (IRB-AAHRPP standards)

Based on the analysis of the theme evolution, we predict that the future AI policy of universities will show three major trends:

From prohibition to guidance: as shown in the guidelines of East China Normal University, the new policy emphasizes "reasonable use" rather than simple prohibition, reflecting a shift in the concept of governance. Technology-enhanced governance: many universities have begun to explore new technological means such as blockchain authentication and AI detection.

Technology-enhanced governance: many schools have started to explore new technological tools such as blockchain depository and AI testing, such as North China Electric Power University's introduction of the "AIGC Testing Service System" .

Integration of literacy education: Xiamen University and other universities have opened courses on AI application and practice, incorporating training on the use of the tools into the formal curriculum.

Author Contributions

C.F. and R.F.: conceptualization, methodology, software; C.F.: data curation, writing—original draft preparation; L.X., F.W. and Y.W.: visualization, investigation; L.X., F.W., Y.W. and S.H.: supervision; C.F. and R.F.: software, validation; C.F. and R.F.: writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

The study did not involve human participants, animal subjects, or sensitive data. Therefore, institutional review board (IRB) approval was not required in accordance with the guidelines of the Declaration of Helsinki.

Informed Consent Statement

Not applicable. The study did not involve human participants, and thus no informed consent was required.

Data Availability Statement

Data supporting the findings of this study are derived from publicly available policy documents of universities included in the analysis (e.g., Fudan University, East China Normal University). The full dataset, including preprocessed text corpora and LDA model outputs, is available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

References

Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, *20*(1), 38. https://doi.org/10.1186/s41239-023-00408-3.

Cotton, D. R., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*(2), 228–239. https://doi.org/10.1080/14703297.2023.2190148.

Eaton, S. E. (2023). Postplagiarism: transdisciplinary ethics and integrity in the age of artificial intelligence and neurotechnology. *International Journal for Educational Integrity*, *19*, 23. https://doi.org/10.1007/s40979-023-00144-1.

Esterhuizen, M. (2025, January 30). *AI adoption in higher education: Bridging the STEM and non-STEM divide*. Phys.org. https://phys.org/news/2025-01-ai-higher-bridging-stem.html.

Hwang, G.-J., & Chen, N.-S. (2023). Editorial position paper: Exploring the potential of generative artificial intelligence in education: Applications, challenges, and future research directions. *Educational Technology & Society*, *26*(2), I–XVIII. https://doi.org/10.30191/ETS.202304_26(2).0014.

Intelligent.com. (2023, January 23). *Nearly 1 in 3 college students have used ChatGPT on written assignments*. https://www.intelligent.com/nearly-1-in-3-college-students-have-used-chatgpt-on-written-assignments.

Kamilia, R. (2024). Exploring the Use of ChatGPT in Teaching Arabic Case Endings: Effectiveness, Challenges and Recommendations. *Journal of Educational Technology and Innovation*, *6*(4), 1–20. https://doi.org/10.61414/jeti.v6i4.198.

Rodgers, J. (2024, January 16). *Oxford releases new guidance on AI use for students*. Cherwell. https://www.cherwell.org/2024/01/16/oxford-releases-new-guidance-on-ai-use-for-students.

Schisgall, E. J. (2023, September 1). *Harvard releases guidance for AI use in classrooms*. The Harvard Crimson. https://www.thecrimson.com/article/2023/9/1/fas-ai-guidance.

Smith, S., Tate, M., Freeman, K., Walsh, A., Ballsun-Stanton, B., Hooper, M., & Lane, M. (2024). A University Framework for the Responsible use of Generative AI in Research. *arXiv*. https://doi.org/10.48550/arXiv.2404.19244.

Thesify. (2024). *Generative AI policies at the world's top universities*. Thesify.ai. https://www.thesify.ai/blog/gen-ai-policies-of-the-worlds-top-universities.

Twabu, K. (2024). Unveiling power dynamics in AI-enabled education: A Foucauldian perspective. *Annals of the University of Craiova, Series Psychology*, *Pedagogy*, *46*(2), 162–176.

University of North Texas. (2023). *Using AI in the higher education classroom*. Digital Strategy. https://digitalstrategy.unt.edu/clear/teaching-resources/theory-practice/using-ai-in-higher-education-classroom.html.

Wenhui Daily. (2024, December 7). *Fudan University responds to the controversy over the "strictest AI regulations"*. ScienceNet. https://news.sciencenet.cn/htmlnews/2024/12/535408.shtm.

Wong, E. (2023, April 14). *UNESCO: ChatGPT, AI, higher education*. Global MOOC and Online Education Alliance. https://www.mooc.global/enoch/unesco-chatgpt-ai-higher-education.

Wong, E. (2024, September 1). *Beyond the horizon: Tsinghua's guide to AI-empowered universities*. CGTN. https://news.cgtn.com/news/2024-09-01/Beyond-the-horizon-Tsinghua-s-guide-to-AI-empowered-universities-1wy0wLfb6BW/index.html.

Wu, C., Zhang, H., & Carroll, J. M. (2024). AI governance in higher education: Case studies of Big Ten universities. *arXiv*. https://doi.org/10.48550/arXiv.2409.02017.

Zhai, X. (2022). ChatGPT User Experience: Implications for Education. *SSRN*. https://doi.org/10.2139/ssrn.4312418.