

Governing Generative Artificial Intelligence: Institutional Policies and Guidelines at America's Flagship Universities

Educational Policy

1–36

© The Author(s) 2026

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/08959048261450995

journals.sagepub.com/home/epx

Jason LaFrance¹ 

Abstract

This study presents a comparative policy mapping analysis of publicly available generative AI policies and guidance across America's flagship public universities, based on documents collected in 2025 and analyzed through thematic coding informed by Clark's Triangle of Coordination. Findings indicate a distributed governance architecture in which instructional decisions are typically delegated to the course level and supported through instructor-facing resources, while institution-level guidance emphasizes data protection, approved tools, and risk management. Although institutions vary in posture, codification, and the depth of guidance provided, the findings point to structural convergence around this shared governance model, suggesting that flagship universities are responding to generative AI in ways that reinforce established distributions of authority within public higher education systems.

Keywords

artificial intelligence (AI), generative AI (GAI), higher education policy, academic integrity, institutional governance, flagship universities, AI ethics and compliance, data privacy in education, technology adoption in academia

¹Appalachian State University, Boone, NC, USA

Corresponding Author:

Jason LaFrance, Department of Educational Leadership, Media, Career Studies, and Leadership Development, Appalachian State University, 151 College Street, Boone, NC 28608-0001, USA.
Email: lafranceja@appstate.edu

Introduction

The emergence of generative artificial intelligence (GAI) tools represents a transformative shift in the landscape of higher education. These tools, powered by large language models and deep learning algorithms, have introduced new possibilities for personalized learning, assessment redesign, and administrative efficiency (LaFrance, 2025a; Michel-Villarreal et al., 2023; Wang et al., 2024). At the same time, they raise critical questions about academic integrity, data privacy, authorship, and the ethical use of algorithmic decision-making (Al-Zahrani & Alasmari, 2024; Cotton et al., 2024). While GAI is a subset of a larger category of artificial intelligence tools, policies have mostly been written in response to the emergence of GAI. Because institutional policies and guidance often use the terms artificial intelligence and generative artificial intelligence inconsistently, this study treats them as overlapping in practice while recognizing that some governance artifacts may reflect broader institutional AI or data compliance structures rather than generative AI-specific policy alone. Institutional policies and guidelines have become key instruments in shaping appropriate use while ensuring compliance as universities grapple with the potential and perils of these technologies (Dabis & Csáki, 2024; McDonald et al., 2025).

Flagship public universities, often viewed as policy leaders within their respective states and at the national level, play a pivotal role in setting the tone for AI policy in particular and artificial intelligence (AI) governance as a whole in higher education (Spivakovsky et al., 2023). While early accounts of institutional responses emphasize variation (with some institutions embracing AI as a pedagogical and research tool and others adopting more cautious or restrictive approaches; Xiao et al., 2023), these differences raise a broader question about whether institutions are developing distinct governance strategies or converging around shared approaches. In this context, higher education leaders face the dual challenge of fostering innovation while navigating a complex web of legal, ethical, and cultural considerations (Gianni et al., 2022; Slimi & Carballido, 2023).

This study contributes to the emerging body of research on AI in education by analyzing how flagship universities frame and implement institutional AI policies and guidelines. Flagship universities, which serve as leading public institutions in their states, are often the first to articulate policy positions on emerging technologies, and their responses can shape system-wide practices and influence national discourse. Drawing on comparative policy analysis of institutional documents, this study examines key domains such as academic integrity, governance structures, data security, and stakeholder engagement.

While prior research has often emphasized variation in institutional responses to generative AI, the findings of this study point toward a different pattern. Rather than occupying a continuum between enthusiasm and caution, flagship public universities exhibit structural convergence around a shared governance model. In this study, convergence refers to a shared governance architecture, specifically the recurring combination of course-level instructional discretion with centralized oversight of data protection, approved tools, and risk management, even when institutions vary in posture, codification, and scope. This convergence is analytically significant. It suggests that generative AI governance is not producing widely divergent institutional strategies but instead reinforcing established distributions of authority within higher education systems. Flagship universities, operating under similar regulatory, reputational, and organizational pressures, appear to be arriving at similar governance arrangements. This pattern is also consistent with institutional isomorphism, in which organizations facing common pressures and uncertainty adopt similar structures over time (DiMaggio & Powell, 1983). At the same time, this convergence can be understood through Clark's (1983) Triangle of Coordination, in which academic authority governs instructional decisions while state and administrative authority are concentrated in compliance and risk domains.

This study extends prior work by examining how governance patterns are structured within a defined institutional context and by interpreting the observed convergence through the lens of authority distribution. While earlier research has mapped institutional policies across mixed institutional types, including private and public universities with differing governance arrangements, this study focuses on a more defined institutional population. By examining flagship public universities within state systems, the analysis captures governance patterns that emerge under shared conditions of public accountability, regulatory oversight, and system-level influence. In addition, it helps explain why institutions arrive at similar governance arrangements while still varying in posture, codification, and scope.

Background and Context

The rapid advancement of generative AI technologies has created both opportunity and uncertainty within higher education. Since the release of publicly accessible tools such as ChatGPT in late 2022, universities have faced increasing pressure to respond to their pedagogical, administrative, and ethical implications.

While generative AI offers the potential to transform learning and institutional efficiency, its adoption has outpaced the development of formal

governance structures, raising questions about transparency, academic integrity, data protection, and equitable access (McDonald et al., 2025; Popenici & Kerr, 2017). Institutional leaders must therefore balance innovation with legal and ethical responsibilities, including privacy, compliance, and responsible use (Gianni et al., 2022; International Association of Privacy Professionals (IAPP) & FTI Consulting, 2024).

Recent scholarship has begun to examine how higher education institutions are responding to generative AI. For example, An et al. (2025) investigated institutional guidelines and policies across a mix of private and public universities identifying emerging practices in teaching, learning, research, and administration. While this work provides an important foundation, it remains largely descriptive and focuses on a heterogeneous set of elite institutions. The present study extends this line of inquiry by examining a more defined institutional population, America's 50 flagship public universities, which operate within distinct contexts of state governance, public accountability, and system-level influence. Rather than focusing solely on variation across institutions, this study examines whether a shared governance pattern is emerging, particularly in how instructional decision-making and institutional oversight are structured. Drawing on Clark's (1983) Triangle of Coordination, the analysis situates this pattern within broader dynamics of higher education governance, highlighting how institutions balance academic authority, administrative control, and external pressures in responding to generative AI.

Purpose and Significance

The purpose of this exploratory study was to examine how flagship public universities in the United States are providing guidelines and implementing institutional policies governing the use of artificial intelligence. Through a comparative policy analysis of publicly available institutional websites and documents this research explores how institutions navigate tensions between innovation, academic integrity, and regulatory compliance in this evolving technological landscape.

This work is significant because it addresses a critical gap in understanding how higher education institutions are responding to the rapid emergence of AI. Institutional policies and guidelines play a pivotal role in shaping equitable and responsible practices as these tools increasingly influence teaching, research, assessment, and administration. By analyzing flagship institutions that often serve as models for state systems and national peers this study contributes to broader conversations on ethical governance, institutional capacity, and policy enactment.

Literature Review

AI in Higher Education: Emerging Landscape

Generative artificial intelligence (GAI) has rapidly emerged as both a disruptor and an accelerator in higher education. Unlike earlier digital tools, GAI produces original text, code, data, images, and video, creating new opportunities to personalize learning, enhance research productivity, and redesign assessment practices (Batista et al., 2024; LaFrance, 2025b; Michel-Villarreal et al., 2023). At the same time it raises profound ethical and pedagogical dilemmas regarding how teaching, learning, and knowledge creation are conceptualized within the academy (Luckin, 2018). Institutional responses to these challenges have been inconsistent (ranging from cautious adoption to outright bans; Wang et al., 2024). Clearly, effective integration requires governance structures, faculty development, and useful frameworks that balance innovation with risk management (Chan, 2023; Southworth et al., 2023). Despite the recognized need for structures and frameworks, sector-level evidence suggests a persistent gap between recognition and institutional readiness. In the 2025 EDUCAUSE AI Landscape Study, 39% of respondents reported that their institution has AI-related acceptable use policies, up from 23% in the 2024 study, indicating movement toward formalization but continued unevenness across the sector (Robert, 2024; Robert & McCormack, 2025). The same study signals ongoing risk-governance challenges. Only 9% of respondents indicated that existing cybersecurity and privacy policies sufficiently address AI-related risks, while 44% reported their institution is implementing or improving cybersecurity policies or guidelines to address AI products (Robert & McCormack, 2025). Respondents also perceived that students use AI more than faculty, reinforcing the urgency of clear student-facing expectations and implementable instructional guidance (Robert & McCormack, 2025). Together, these patterns highlight that adoption is accelerating faster than coherent governance, leaving institutions to bridge policy and practice through evolving guidance and support structures. Ghimire and Edwards (2024) further note that where policies exist, they often fail to address student privacy and algorithmic transparency.

Among the areas where these policy gaps are most visible and contested is academic integrity. Complicated questions related to plagiarism, authorship, and cheating remain an unsolved conundrum. One reason is highlighted by Cotton et al. (2024) as they emphasize the difficulty of distinguishing between legitimate AI-supported learning and academic dishonesty. This aligns with the challenges Xiao et al. (2023) highlight related to uneven enforcement and confusion within and across institutions. Wang et al. (2024) report that most universities have issued general guidance rather than enforceable policies,

often leaving decisions to individual instructors. This decentralized approach creates gaps in implementation, especially given the unreliability of AI detection tools (McDonald et al., 2025). As Dabis and Csáki (2024) argue, academic integrity in the age of GAI requires the development of new norms, pedagogies, and enforcement structures. Beyond academic integrity, these debates point to a broader set of concerns that extend to the core of ethical AI use in higher education. Institutions must also contend with issues of data privacy, algorithmic bias, and responsible governance to ensure that GAI adoption supports equitable and ethical educational practices. Emerging empirical research has begun to systematically examine how higher education institutions are responding to generative AI at the policy level. For example, An et al. (2025) analyzed institutional guidelines across mixed institutional types, identifying patterns in teaching, learning, research, and administrative guidance. Similarly, McDonald et al. (2025) examined policies across U.S. institutions, finding that many universities rely on guidance, sample language, and recommended practices rather than enforceable instructional rules. Xiao et al. (2023) and Wang et al. (2024) likewise document substantial variability in institutional responses, with many universities delegating decision-making to individual instructors. Complementing this work, Ganguly et al. (2025) examined guidance for research use across 30 R1 institutions and found that institutions place primary responsibility on individual researchers to interpret and comply with evolving expectations related to disclosure, attribution, and ethical use. Dotan et al. (2024) further highlight how traditions of shared governance and academic freedom shape institutional approaches, often placing higher education in tension with centralized, top-down governance models common in other sectors.

While these studies provide important early insights, they often examine heterogeneous institutional samples or focus on general trends across the sector. Less attention has been given to how a specific institutional population, such as flagship public universities, structures governance in relation to state systems, public accountability, and system-level influence. By focusing on this population, the present study extends existing work by examining not only what institutions are doing, but how governance authority is distributed across instructional and institutional domains.

Development and Governance: Legal and Ethical Frameworks Shaping Institutional Policy

Policy development in higher education is shaped not only by formal legal frameworks but also by how institutions interpret and operationalize those requirements in practice. Universities must navigate a complex regulatory

environment that includes federal, state, and international data privacy and security laws, often applying statutes that were not designed with generative AI in mind. As a result, institutional policy development involves interpretation, risk management, and cross-functional coordination rather than straightforward compliance. These challenges are compounded by longstanding tensions in higher education governance, including balancing academic freedom with institutional accountability and aligning the interests of multiple stakeholders such as administrators, faculty, and students (Ren & Li, 2013; Rizvi & Lingard, 2010). Within this context, generative AI introduces additional uncertainty, as institutions must determine how emerging tools fit within existing policy domains while maintaining flexibility in teaching and learning environments.

In the United States, institutions operate within a layered regulatory landscape that shapes how AI-related policies are developed and enacted. Federal laws such as the Family Educational Rights and Privacy Act (FERPA) and the Health Insurance Portability and Accountability Act (HIPAA), along with international frameworks such as the General Data Protection Regulation (GDPR), establish expectations for data privacy, consent, and security. Additional statutes, including the Gramm-Leach-Bliley Act (GLBA) and the Federal Information Security Management Act (FISMA), further define institutional responsibilities for financial and administrative data. These frameworks do not prescribe specific approaches to generative AI, but they create boundary conditions that institutions must interpret when developing policies and guidance. A central challenge is that many of these legal frameworks were developed prior to the emergence of generative AI and do not explicitly address its capabilities or risks. As a result, institutional leaders must interpret how existing requirements apply to new forms of data use, automated content generation, and algorithmic decision-making. This interpretive work introduces uncertainty and requires institutions to anticipate potential risks related to privacy, bias, and compliance without consistent or unified regulatory guidance. Consequently, policy development becomes an ongoing, adaptive process that relies on collaboration among legal, academic, and technical stakeholders (Airaj, 2024; Dignum, 2020; Gianni et al., 2022; IAPP & FTI Consulting, 2024; Slimi & Carballido, 2023).

These conditions help explain why institutional responses to generative AI frequently emphasize guidance, risk mitigation, and adaptable governance structures rather than fixed, comprehensive policies. As reflected in sector-level evidence, institutions are expanding AI-related policies and practices, but governance remains uneven and evolving across domains such as teaching and learning, data privacy, and institutional operations (Robert, 2024; Robert & McCormack, 2025). These dynamics highlight that generative AI

governance in higher education is not simply a matter of compliance, but of how authority is distributed and exercised across institutional domains. Understanding these patterns requires a framework that can account for the interaction of regulatory pressures, institutional priorities, and academic decision-making.

Theoretical Framework

This study is guided by Clark's (1983) Triangle of Coordination, a foundational framework in higher education policy that explains how authority is distributed across three domains: state authority, market forces, and academic oligarchy (often described in contemporary literature as academic authority). State authority reflects regulatory, legal, and compliance-driven influences on institutional decision-making. Market forces represent pressures related to competition, innovation, and institutional positioning. Academic authority refers to the influence of faculty expertise, disciplinary norms, and shared governance structures on teaching and academic practice. In the context of generative AI governance, these domains are reflected in distinct but interconnected ways. State and administrative authority are reflected in institutional policies related to data protection, privacy, and approved technologies. Market pressures are visible in institutional efforts to signal innovation, support AI literacy, and maintain competitiveness in a rapidly evolving technological landscape. Academic authority is evident in the delegation of instructional decision-making to faculty at the course level. Clark's framework provides a lens for examining how institutions balance these influences when developing governance approaches to generative AI, with the relative strength of each domain shaping how governance is enacted in practice.

Methodology

This exploratory study employs a comparative policy analysis design utilizing qualitative content analysis to examine the implementation of AI policies at flagship universities in the United States. The comparative approach allows for identifying similarities, differences, and unique practices across institutions. Qualitative content analysis enables a deeper understanding of the themes, patterns, and contextual nuances within institutional policy and guidance materials (Bowen, 2009; Merriam, 2009; Patton, 2015). Together, these methodologies provide a robust framework for systematically analyzing and interpreting the data to generate meaningful insights into policy frameworks and their implications for higher education governance (Creswell, 2021; Yin, 2018).

Research Question

How do flagship public universities in the United States govern the use of generative AI and what does that mean for instructors and students?

Institutional Sample

This study analyzes institutional AI policies at 50 flagship public universities in the United States, selected through purposeful sampling to reflect a broad cross-section of geographic regions, institutional contexts, and enrollment characteristics (see Table A1). Flagship universities were chosen because they are the leading public institutions in their respective states, typically characterized by high research activity, substantial public funding, and a statewide mandate for educational leadership. In addition to state appropriations and tuition revenue, these institutions generate significant external research funding and frequently influence policy development and best practices across state systems and national contexts. Their resource capacity also positions them to develop institution-wide guidance, invest in emerging technologies, and support faculty development in areas such as generative AI. The sample offers meaningful diversity across several dimensions. Geographically, the universities represent all 50 states, spanning coastal, midwestern, southern, and mountain regions. According to NCES (2023), most flagship institutions are situated in city-based contexts, with 74% ($n=37$) classified as urban, 14% ($n=7$) located in suburban areas, and 12% ($n=6$) in town or rural settings (see Tables A1 and B1). Of the 50 flagship universities analyzed, 45 are classified as R1 (Doctoral Universities-Very High Research Activity) and five are classified as R2 (Doctoral Universities-High Research Activity): the University of Alaska Fairbanks, University of Idaho, University of Rhode Island, University of Vermont, and the University of Wyoming. In terms of student enrollment, undergraduate populations range from 5,638 students at the University of Alaska Fairbanks to 45,728 at Ohio State University, while graduate enrollment ranges from 1,002 at the University of Alaska Fairbanks to 19,890 at the University of Florida. The median undergraduate enrollment across flagship universities is approximately 24,685 and the median graduate enrollment is 8,073. These characteristics underscore that flagship institutions are large, complex organizations with diverse academic and administrative units, making them analytically appropriate sites for examining institutional governance structures.

Focusing on flagship public universities, this study is informed by the expectation that governance approaches to generative AI may differ from those of private or non-flagship institutions. Public flagships operate within

systems of state oversight and public accountability, which may lead to stronger emphasis on compliance, data governance, and risk management. At the same time, their size, research intensity, and resource capacity position them to develop institution-wide guidance and support structures more readily than smaller or less resourced institutions. Compared to private institutions, which may have greater flexibility in policy design, flagship universities may exhibit more standardized or system-aligned approaches. Beyond these expectations, flagship universities occupy a distinctive position within higher education systems. As state-designated institutions with broad access mandates and substantial public visibility, they often function as policy leaders within their respective states and influence system-wide practices. Examining generative AI governance within this population therefore provides insight into how public institutions balance innovation, compliance, and accountability, offering a perspective that may not be fully captured in studies focused on private or mixed institutional samples.

Data Collection Procedures

Policy documents were collected in January 2025 and updated in August 2025 through systematic website searches. Site-specific queries were conducted using each institution's name in combination with key terms such as "generative AI policy," "artificial intelligence policy," "academic integrity policy," "artificial intelligence data protection," "AI guidelines," "AI syllabus statement," and "AI use policy." In addition, internal links within institutional websites were followed to locate related governance materials. Beyond explicit AI policies and guidance pages, additional documentation was reviewed, including academic integrity policies, student handbooks, and student codes of conduct. During document review, explicit references to artificial intelligence were identified (e.g., "artificial intelligence," "AI," "generative AI," "large language model," "automated text generation," or equivalent phrasing).

During initial collection, policy text was captured through direct copy-and-paste into Word files for archival purposes and, where feasible, imported into Dedoose. As the corpus expanded and institutional guidance proved highly distributed across interconnected webpages, the collection strategy shifted to structured research notes recorded in a standardized metadata log. This log documented page titles, URLs, artifact type (e.g., academic integrity, IT or data governance, teaching and learning guidance), and analytic memos summarizing substantive expectations. This approach supported consistent cross-institution comparison while accommodating the web-native, multi-page structure of institutional AI governance materials.

Given the distributed and web-native nature of institutional AI governance, the analytic corpus is best understood as a collection of policy artifacts rather than discrete, stand-alone policy documents. Artifacts were categorized into two broad groups: primary governance artifacts and supporting resources. Primary artifacts included institutional policies and guidance that articulated formal expectations or rules related to AI use, such as academic integrity codes, IT or data governance guidance, approved tool lists, and institution-level AI guidance pages. Supporting artifacts included instructional and informational resources intended to facilitate implementation, such as teaching center webpages, sample syllabus statements, frequently asked questions, and student-facing AI resource pages. Across institutions, this approach yielded an estimated four primary artifacts and four supporting artifacts per campus, resulting in an analytic corpus of approximately 400 artifacts. Categorization was used to support consistent cross-institution comparisons rather than to imply equivalent policy weight across artifact types.

Coding and Thematic Analysis Process

The study employed thematic analysis to systematically analyze the collected data, ensuring a nuanced understanding of the content and context of AI policies (Braun & Clarke, 2006; Saldaña, 2021). The first step involved an initial review of the archived materials and research notes to familiarize the researcher with institutional scope and context and to identify recurring policy topics. Open coding was then applied to identify salient governance features, including ethics, privacy, assessment expectations, disclosure norms, and enforcement mechanisms (Saldaña, 2021). Codes were iteratively refined and then grouped into overarching themes. Finally, themes were compared across universities to identify cross-institution patterns, points of divergence, and distinctive practices (Creswell & Poth, 2018). Descriptive frequencies (n and %) were calculated from coded indicators to contextualize patterns across institutions and to anchor narrative claims in the Findings section. Key indicators and decision rules are summarized in Table C1.

To enhance analytic rigor, generative AI tools were used as an auxiliary analytic aid during interpretation and write-up. Specifically, ChatGPT (version 5.1) was used to support sense-making by generating alternative interpretations of emergent themes, suggesting potential blind spots, and helping refine category boundaries through iterative prompt-based inquiry. Because the corpus consisted of publicly available institutional policy and guidance materials, excerpts and summaries were drawn directly from the collected webpages and documents, and the coded spreadsheet was used as the primary structured record of results. Data-sharing was disabled in ChatGPT settings

during use. Outputs were treated as provisional suggestions rather than evidence. They were reviewed against the original sources and the researcher's codebook and were used to prompt re-checking and clarification rather than to replace human judgment. This approach is consistent with recent recommendations for cautious, transparent use of AI tools in qualitative research (Barany et al., 2024; Marshall & Naff, 2024; Morgan, 2023; Theelen et al., 2024).

Policy orientation classification was treated as a focused, downstream coding step within this thematic process. Each institution was coded for its dominant discourse pattern based on the overall tone of its governance posture across primary artifacts. Policies emphasizing prohibitive or compliance-oriented language (e.g., "unauthorized use," "subject to disciplinary action") were categorized as Restrictive. Policies frequently using conditional or guidance-oriented terms (e.g., "with instructor permission," "responsible use") were categorized as Balanced. Policies emphasizing innovation, ethics, or AI literacy (e.g., "integration into teaching and learning," "AI literacy," "bias mitigation") were categorized as Transformative. Where institutions published multiple artifacts, this classification reflected the prevailing posture across the institution's primary governance materials rather than isolated statements on a single page.

Ethical Considerations

No human subjects were involved in this study. All data analyzed were publicly available institutional documents. Ethical considerations focused on transparency and responsibility in using GAI to support analysis. The role of ChatGPT 5.1 was clearly delineated as a tool for triangulation and exploration rather than a substitute for human judgment. Following the guidance of Longo (2019) and Marshall and Naff (2024), special care was taken to verify AI-generated outputs, document all prompt interactions, and ensure that interpretive authority remained with the human researcher.

Findings

Guided by the question of how flagship public universities govern generative artificial intelligence through guidance and policy, the analysis points to a consistent institutional choice to rely on guidance rather than a single binding instructional policy, while establishing more formal guardrails in the domains of data protection and tool governance. Across the sample, campuses set baseline expectations for privacy and data classification, publish vetted or approved tools, and rely on instructor discretion for course level decisions.

This configuration reflects a distributed governance architecture in which authority is differentiated across institutional and instructional domains. Findings are organized into four themes: (1) common core elements campuses share, (2) institutional differences, (3) instructor-facing guidance, and (4) student-facing guidance.

Descriptive Overview of Policy Artifacts

Descriptive frequencies are provided to contextualize patterns identified through qualitative coding rather than to support inferential analysis. Across the 50 flagship universities analyzed, course-level discretion over instructional AI use was universal ($n=50$, 100%), with most institutions supplying sample syllabus statements ($n=44$, 88%). Discrete institutional AI guidance from Information Technology (IT departments; e.g., a dedicated AI webpage or formal guidance document) was present at 42 institutions (84%). By contrast, explicit references to AI within academic integrity codes appeared at 17 institutions (34%), with the remainder relying on broader unauthorized-assistance language and parallel AI guidance documents. Data and tool governance guidance was prevalent ($n=46$, 92%), typically through recommended approved-tool lists ($n=39$, 78%) or mandated use of approved tools in defined contexts ($n=7$, 14%). Transparency norms were widely articulated, with 15 institutions (30%) mandating disclosure and 34 (68%) recommending disclosure. Policy posture varied across institutions: balanced or guidance-oriented ($n=32$, 64%), transformative or innovation-oriented ($n=10$, 20%), and restrictive or compliance-oriented ($n=8$, 16%). Taken together, these descriptive patterns indicate a strong baseline consistency in how institutions structure governance, even as specific policy expressions vary across campuses. See Table C1 for the full indicator set.

Common Core Elements. The review of institutional materials indicated that the most visible artifacts of governance were guidance resources rather than consolidated, binding instructional policies. Across institutions, guidance commonly took the form of dedicated AI webpages, teaching-center resources, and model syllabus language that supported local decision-making. Consistent with this approach, all 50 institutions (100%) located instructional decisions about student AI use at the course level, and most provided sample syllabus statements ($n=44$, 88%). In addition, 42 institutions (84%) published standalone AI guidance housed in IT (or adjacent administrative units) rather than within academic integrity codes, reflecting a governance logic in which institutions provide guardrails and resources while instructors establish course-specific boundaries.

Instructional flexibility was typically paired with institution-level guardrails regarding data and tools. Data and tool governance guidance was prevalent ($n=46$, 92%), most commonly through recommended lists of approved tools ($n=39$, 78%) and, less frequently, through mandated use of approved tools in defined contexts ($n=7$, 14%). These guardrails generally emphasized data classification, privacy, and contractual protections. In four institutions (8%), a comparable approved-tools list or procurement guidance was not readily available in public-facing materials at the time of data collection, suggesting variation in how explicitly tool governance is documented. This pairing of decentralized instructional authority with centralized data governance represents a defining feature of the observed governance architecture.

A second recurring element was caution regarding the limits of AI detection tools in academic integrity processes. Faculty were discouraged from relying on AI detection software as the sole basis for misconduct decisions, with several campuses emphasizing corroborating evidence and due process. Within the broader mapping, this functions as an implementation signal: institutions are encouraging educators to prioritize clear expectations, assessment design, and documentation over surveillance-based approaches.

Finally, several campuses explicitly described why they have not adopted a comprehensive, universal instructional AI policy, instead framing governance as an evolving set of resources responsive to rapid technological and regulatory change. For example, the University of Idaho recommends that the institution “create and maintain guidance” across teaching, research, and administration and concludes that “it is not recommended that a specific policy be developed to address AI,” while the University of Kansas similarly urges instructors to articulate clear syllabus expectations as broader institutional approaches continue to develop. Taken together, these examples reinforce a common governance architecture across flagships: course-level discretion for instructional use paired with institution-level guardrails for data and tool governance, alongside caution about relying on AI detection tools as the primary basis for academic integrity enforcement.

Institutional Differences. Even with shared guardrails, campuses differ in three ways that shape governance. The first is locus of decision-making and the degree of institutional scaffolding provided to instructors. While all institutions locate instructional decisions at the course level, they vary in how much pre-packaged language, training, and unit-level guidance they provide. This variation is reflected in broader policy posture: 32 institutions (64%) were coded as balanced or guidance-oriented, 10 (20%) as transformative or innovation-oriented, and 8 (16%) as restrictive or compliance-oriented. These postures signal not only differences in tone but also in how institutions frame

generative AI (as a risk to be managed, a tool to be governed through shared norms, or an opportunity for curricular and institutional innovation). Transformative orientations, for example, are exemplified by Ohio State University's stated commitment to AI fluency for all students in the class of 2029 and the University of Florida's AI Across the Curriculum initiative. In these cases, governance extends beyond risk mitigation to include explicit investments in capacity-building, signaling an expectation that faculty will engage with AI as a pedagogical and professional resource rather than merely a compliance concern.

The second difference is how academic integrity violations are codified. Although institutions consistently emphasize instructor authority and course-level rules, only 17 of the 50 institutions (34%) explicitly reference artificial intelligence in their academic integrity codes. The remaining 33 institutions (66%) rely on broader unauthorized-assistance language under which AI-related violations may be adjudicated, often supplemented by separate guidance documents. This pattern suggests uneven policy revision practices across institutions, with some updating integrity codes directly and others maintaining legacy code language while publishing parallel AI guidance.

The third area of divergence concerns governance beyond the classroom, particularly guidance for research and institutional operations. Research integrity and compliance considerations were addressed in 33 institutions (66%) but were not evident in the public-facing materials of 17 (34%), suggesting that research-related governance is less consistently articulated than teaching-focused guidance. Where present, research guidance commonly emphasizes confidentiality, publication and grant compliance, and restrictions on using public AI tools with sensitive or proprietary data. Several institutions also extend governance to high-consequence administrative uses, outlining limits or consultation requirements for activities such as hiring, discipline, and automated grading. Illustratively, the University of Alabama cautions that researchers "remain responsible" for AI-assisted content and must safeguard sensitive information, University of California–Berkeley flags "highly consequential automated decision making" as requiring additional consultation, the University of Georgia's Graduate School restricts AI use in theses and dissertations absent explicit approval, and the University of Minnesota limits AI use with private data and for generating non-public institutional outputs. Table 1 summarizes areas of convergence and divergence across institutions.

Instructor-Facing Guidance. Instructor-facing guidance emerged as a central implementation feature across institutional materials and was most commonly operationalized through ready-made syllabus statements and pedagogical

Table 1. Convergence and Divergence in Flagship University AI Governance.

Area of Policy	Convergence	Divergence
Data Protection	Most institutions provide data classification guidance and identify approved tools for AI use ($n=46$, 92%).	Approved-tool lists or procurement guidance were not readily available in four institutions ($n=4$, 8%) in public-facing materials at the time of collection.
AI Detection Tools	Widespread caution against using AI detection outputs as the sole basis for misconduct decisions.	Some institutions discourage use entirely; others allow cautious reference as a preliminary flag requiring corroboration.
Academic Integrity	Shared emphasis on instructor authority and course-level rules.	Explicit AI language in integrity codes ($n=17$, 34%) versus reliance on broader unauthorized-assistance language ($n=33$, 66%).
Research & Administration	Research integrity/compliance considerations appear in a majority of institutions ($n=33$, 66%).	Some institutions specify restrictions or consultation requirements for high-consequence uses (e.g., UC Berkeley's guidance on "highly consequential" automated decisions; Georgia Graduate School guidance on theses/dissertations).
Syllabus Guidance	Most institutions provide sample syllabus statements or templates ($n=44$, 88%).	Some provide detailed unit-level frameworks (e.g., Alabama, Georgia), while others provide only general reminders to articulate course expectations.

Note: "Convergence" indicates patterns common across institutions; "Divergence" summarizes meaningful differences in how institutions address the policy area.

prompts. Across the sample, 44 institutions (88%) provide sample syllabus language or templates to support course policy communication, while the remaining six (12%) encourage instructors to articulate expectations without supplying model text. Where templates are provided, they typically offer a spectrum of options from prohibitive to conditional to permissive use, intended to support alignment with course learning outcomes and assessment design.

In many cases, these syllabus resources are paired with design-oriented prompts housed in teaching and learning centers. These prompts encourage instructors to consider how AI use aligns with course objectives and to design assessment strategies that make student thinking visible. Equity and access considerations appear less consistently in the public-facing guidance, explicitly referenced at 19 institutions (38%). Instructor-facing guidance also commonly aligns with institution-wide tool and data guardrails, reinforcing that course-level discretion operates within broader expectations for privacy, approved tools, and responsible use.

Taken together, the implementation logic is consistent across institutions. Faculty are encouraged to establish and communicate a clear stance early, connect that stance to learning goals, and design assessments that foreground process, attribution, and student responsibility for submitted work. This approach is reinforced by transparency norms. Disclosure is mandated in 15 institutions (30%) and recommended in 34 (68%), with only one institution not addressing disclosure explicitly in the coded guidance. A consolidated summary of stakeholder-facing guidance patterns is provided in Table D1.

Student-Facing Guidance. Student-facing recommendations were commonly articulated through academic integrity policies, honor codes, and AI resource pages. These materials emphasize two recurring expectations: clarification and transparency. Because acceptable AI use is typically determined at the course level, students are routinely advised to review course policies and seek clarification from instructors before using AI tools on coursework. This “ask first” norm is reinforced by institutional messaging that AI expectations may vary across courses, instructors, and programs.

A second expectation is that permitted AI use should be disclosed. Across the sample, 15 institutions (30%) mandate disclosure of AI use in student work, while 34 institutions (68%) recommend disclosure as a best practice or leave documentation requirements to instructor discretion. One institution did not address disclosure explicitly in the coded materials. Where disclosure norms are articulated, they are often paired with reminders that students remain responsible for the accuracy of facts, ideas, and citations. This was true even when AI tools assist with drafting or revision. Student-facing guidance also frequently includes privacy cautions. For example, institutions

warn against inputting sensitive or protected information into public tools and, in some cases, restrict use of private or restricted data to approved environments. Taken together, these messages position clarification and disclosure as operational norms in the emerging governance landscape. See Table D1 for a cross-audience summary of guidance elements.

Discussion

This study examined how America's flagship public universities govern generative artificial intelligence (GAI) and what these approaches mean for instructors and students. Across the 50 institutions analyzed, the findings indicate a consistent governance pattern. Rather than adopting prescriptive, institution-wide instructional policies, flagships have largely relied on guidance documents, data and tool guardrails, and course-level discretion. This pattern aligns with prior research suggesting that many institutions have responded to generative AI through general guidance and resources while leaving key instructional decisions to instructors and academic units (Wang et al., 2024; Xiao et al., 2023). At the same time, the prevalence of data protection and tool governance guidance reflects a parallel emphasis on compliance, privacy, and risk management, which scholars have identified as central tensions in early institutional AI governance efforts (Dabis & Csáki, 2024; Gianni et al., 2022). These findings indicate that, rather than developing highly divergent approaches, institutions are converging around a shared governance architecture characterized by decentralized instructional decision-making and centralized oversight of data protection, approved tools, and risk management. This pattern suggests that generative AI is not fundamentally disrupting institutional governance structures, but instead is being integrated in ways that reinforce existing distributions of authority within higher education systems.

Interpreted in relation to institutional type, the findings suggest that this is not simply a higher education AI governance story, but a public flagship governance story. The convergence observed across institutions appears to reflect the shared conditions of public accountability, state oversight, and system-level leadership that distinguish flagship public universities from more autonomous private institutions.

Distributed Governance and Institutional Caution

The reliance on guidance rather than binding instructional policy reflects a cautious and pragmatic institutional response to a rapidly evolving technology. In Clark's (1983) terms, this governance architecture reflects a

configuration in which academic authority is privileged in the instructional domain, enabling disciplinary and pedagogical variation at the course level. At the same time, the strong presence of institution-level tool vetting and data classification guidance indicates that state authority is expressed through institutional policies related to compliance, data protection, and legal risk, while administrative structures operationalize these requirements through tool vetting and guidance (Gianni et al., 2022). Market pressures are also evident, as institutions seek to balance innovation with reputational and competitive considerations in their approach to generative AI. Within Clark's framework, these patterns reflect a distribution of authority in which academic authority governs instructional decisions, state authority is concentrated in compliance and risk domains, and market forces shape institutional positioning around innovation and competitiveness.

Taken together, these patterns point to a consistent governance model across institutions rather than substantial divergence in approach. Flagship universities have largely converged on a structure that combines decentralized instructional decision-making with centralized oversight of data, tools, and risk. This model allows institutions to remain adaptable in the face of rapid technological change while maintaining control over areas of legal and operational vulnerability, reinforcing established distributions of authority within higher education systems. These findings also reflect the distinctive position of public flagship universities within state systems, where governance must balance innovation with public accountability, regulatory compliance, and system-level coherence. As policy leaders, their convergence around a distributed governance model may signal the emergence of a sector-wide approach that extends beyond individual institutional contexts.

For institutional leaders, this approach offers several advantages. It avoids overregulation, reduces the risk of rapid policy obsolescence, and respects disciplinary differences in how generative AI may be used in teaching and learning. At the same time, this model mirrors uneven institutional readiness reported in sector-level evidence, where policy infrastructure and cybersecurity and privacy preparation for AI remain incomplete even as use expands (Robert & McCormack, 2025). As a result, the reliance on distributed governance shifts significant interpretive and implementation responsibility to faculty, shifts significant interpretive and implementation responsibility downward, with important implications for workload, consistency, and student experience.

Faculty as Determiners of Classroom AI Policy. A central implication of these findings is that institutions have effectively delegated rule-setting authority for instructional AI use to faculty. Instructors function as the primary determiners

of classroom AI policy by establishing and enforcing expectations through course policies, classroom norms, assignment instructions, learning management system communications, and ongoing guidance throughout the term. This observation is consistent with prior scholarship noting that institutional responses frequently rely on instructor discretion and localized decision-making rather than uniform instructional policies (Wang et al., 2024). The widespread provision of sample syllabus statements and teaching center resources further signals that institutions expect faculty to translate broad guidance into enforceable classroom practice, a pattern that is also consistent with calls for assessment redesign and pedagogical frameworks that situate AI use in learning outcomes rather than blanket prohibitions (Chan, 2023; Southworth et al., 2023). Viewed through Clark's (1983) Triangle of Coordination, this pattern reflects the continued dominance of academic authority in the instructional domain. Even in the context of a disruptive technology, institutions largely preserve faculty control over course-level rules rather than centralizing instructional decision-making.

This distributed model aligns with norms of academic self-governance, but it also carries costs. Faculty vary in their familiarity with AI tools, comfort with experimentation, and capacity to redesign assessment practices at scale. As Robert and McCormack (2025) suggest, uneven institutional readiness can amplify uneven implementation, even when guidance exists. In addition, the widespread caution about AI detection tools intensifies faculty responsibility. Research has highlighted both the practical and ethical limitations of AI detection in misconduct processes, including the difficulty of distinguishing legitimate learning support from dishonest substitution and the risks of overreliance on imperfect detection outputs (McDonald et al., 2025). In that context, institutions' emphasis on assessment design, transparency norms, and documentation aligns with the literature, but it also reinforces that faculty are carrying the primary burden of policy determination, communication, and enforcement in the instructional domain (Cotton et al., 2024).

Equity Implications of Decentralized Governance. Although equity and access were not consistently foregrounded in institutional policy artifacts, the governance patterns identified in this study have implications for how expectations are experienced by students. This interpretation is supported by the findings that instructional decision-making was universally located at the course level ($n=50$, 100%) and that explicit equity and access language appeared in only 38% of institutional guidance, indicating that expectations may vary while equity considerations are not consistently embedded across institutions. Because instructional expectations are determined at the course level and vary across instructors, students may encounter inconsistent norms

for acceptable AI use across courses and programs. Prior work has noted that uneven institutional guidance and instructor discretion can contribute to confusion and variability in student experience (Xiao et al., 2023). The readiness gap reported by institutional leaders further suggests that decentralized approaches may contribute to variation in how supports, expectations, and enforcement practices are implemented (Robert & McCormack, 2025). From a theoretical standpoint, this is also notable because equity does not appear as a distinct domain within Clark's Triangle of Coordination. The relative absence of explicit equity language in governance artifacts may therefore reflect not only uneven institutional attention, but also the fact that equity considerations are not strongly embedded in the authority structures most visibly shaping early generative AI governance.

From an equity perspective, the widespread "ask first" norm is especially important to interpret. While asking for clarification is reasonable, it also places responsibility on students to navigate differences across courses. Students who are more familiar with institutional expectations or more comfortable seeking clarification may be better positioned to interpret and respond to these expectations. Conversely, students with less experience navigating institutional norms may face greater uncertainty in how policies are understood and applied. In this sense, decentralized governance may contribute to variation in student experience even when institutional intent is to preserve flexibility and academic freedom. A policy implication for institutional leaders is that equity considerations should be integrated into baseline expectations, communication practices, and accessible supports rather than addressed solely through general statements of principle (Gianni et al., 2022; Robert & McCormack, 2025).

Policy Postures and Institutional Signaling. The variation in institutional posture, balanced, transformative, or restrictive, illustrates that governance choices also function as signals. Balanced approaches emphasize conditional use and responsible experimentation. Transformative approaches frame AI as an opportunity for innovation and literacy. Restrictive approaches emphasize compliance and risk mitigation. These postures influence how faculty and students interpret institutional priorities, and they likely shape where campuses invest in professional learning and infrastructure (Chan, 2023; Southworth et al., 2023). The prevalence of balanced approaches suggests that most flagships are attempting to navigate a middle path between innovation and control, which is consistent with the broader literature describing institutional experimentation and a lack of consensus about how AI should be integrated into academic work (Dabis & Csáki, 2024; Wang et al., 2024). Within Clark's framework, these postures can also be read as responses to market

forces, particularly pressures related to institutional positioning, innovation signaling, and competitive reputation. In that sense, differences in posture reflect not only governance tone but also how institutions present themselves in relation to a rapidly evolving technological environment.

For institutional leaders, posture becomes consequential when it is misaligned across governance domains. For example, a permissive instructional message paired with restrictive tool and data guidance can generate uncertainty about what “responsible use” means in practice. Conversely, restrictive classroom messaging paired with innovation-forward institutional rhetoric can create mixed signals that increase both student confusion and faculty workload. Coherence across instructional guidance, integrity frameworks, and tool governance is therefore a leadership-level governance task, not merely a communication preference.

Implications for Institutional Leadership. Taken together, the findings suggest that flagship universities are governing GAI through a deliberately distributed model that emphasizes guidance, faculty rule-setting in the classroom, and centralized guardrails for data and tools. This arrangement is understandable given uncertainty, rapid technological change, and documented limits in detection and enforcement mechanisms (Cotton et al., 2024; McDonald et al., 2025). However, it also creates predictable risks: fragmented practice, uneven student experience, and equity concerns when expectations vary widely and students must navigate ambiguity across courses (Robert & McCormack, 2025; Xiao et al., 2023). Interpreted through Clark’s Triangle of Coordination, this distributed model reflects a stable allocation of authority across domains rather than a temporary policy gap. Academic authority remains concentrated in instructional decision-making, while state and administrative authority are most visible in compliance, privacy, and risk management.

For institutional leaders, the policy challenge is not whether to centralize or decentralize. It is how to coordinate across levels so that flexibility is principled rather than inconsistent. Establishing shared baselines for disclosure, attribution, and student responsibility can reduce confusion while preserving pedagogical autonomy. Strengthening professional learning and cross-unit coordination can also reduce the burden placed on individual instructors to interpret rapidly changing technologies and translate guidance into enforceable classroom policy (Chan, 2023; Southworth et al., 2023). In this sense, the current moment represents an opportunity for campuses can move from provisional guidance toward more coherent and equity-aware governance without foreclosing innovation.

Limitations and Delimitations

This study is subject to several limitations that should be considered when interpreting the findings. First, the analysis relied exclusively on publicly available institutional policy documents and guidance as of August 2025. As a result, internal policies, informal practices, or subsequent revisions may not be reflected. The findings therefore represent a time-bound snapshot of institutional governance rather than a comprehensive account of ongoing policy development.

Second, the study employed qualitative content analysis, which necessarily involves interpretive judgment. While analytic rigor was supported through systematic coding procedures and the use of generative AI tools to assist with sense-making and category refinement, interpretive authority remained with the researcher. Institutional variability in document structure, terminology (for example, policies versus guidelines), and web-based presentation also posed challenges for direct comparison across cases, despite efforts to standardize coding through a structured analytic framework.

Finally, the analysis focuses on the presence and framing of governance artifacts rather than their enactment. Publicly articulated policies and guidance do not necessarily indicate how rules are implemented, interpreted, or enforced in practice.

The delimitations of the study were intentional and designed to focus the scope of analysis. The sample was limited to flagship public universities because of their visibility and leadership role in shaping higher education policy. Consequently, the findings may not generalize to private institutions, 2-year colleges, or international contexts. The corpus was restricted to finalized, public-facing policy documents and guidance; draft materials, internal communications, and informal faculty practices were excluded. In addition, the coding strategy emphasized themes aligned with the study's research question and theoretical framework, which may have limited attention to other relevant aspects of institutional AI governance.

Recommendations

The findings suggest that flagship universities are building generative AI governance through a distributed instructional model that places classroom rule-setting largely with faculty, while concentrating institution-level authority in data protection, approved tools, and risk management. This approach preserves disciplinary flexibility and can adapt to rapid technological change, but it also increases the likelihood of uneven student experiences and uneven implementation capacity across courses and programs. Given the role of

flagship universities as public, state-affiliated institutions with system-level influence, these dynamics are especially consequential for leaders responsible for coordinating governance across complex institutional environments. The recommendations below are intended for institutional leaders including provosts, Chief Information Officers, academic governance committees, and teaching and learning centers, who must balance academic freedom, equity, compliance, and instructional quality. Because this governance pattern reflects a relatively stable distribution of authority across academic, state-administrative, and market domains, effective institutional responses must work across these domains rather than rely on a single centralized solution.

Establish Campus-Wide Baselines for Instructional AI Use

Institutions can retain course-level discretion while still adopting a limited set of common expectations that apply across the institution. In Clark's terms, this creates a common floor within the instructional domain without displacing faculty authority over course-level implementation. These baselines should be intentionally narrow and implementation-ready, focusing on elements that reduce confusion for students and reduce burden for faculty. At minimum, institutional leaders should clarify expectations for disclosure of AI use (what to disclose and how), attribution norms (how AI assistance should be acknowledged when permitted), and student responsibility (accuracy, originality, and alignment with course requirements). Clear baselines can function as a floor, not a ceiling, so that instructors and academic units can be more restrictive or more permissive in ways that align with learning outcomes and disciplinary norms.

Improve Coherence Across Guidance, Integrity, and Classroom Policy

Many institutions are governing AI through multiple artifacts rather than a single consolidated policy. This is especially important in a governance environment where academic authority and state-administrative authority operate in parallel rather than through a single unified policy channel. That structure can be workable, but only when the message is coherent. Leaders should ensure that AI guidance explicitly cross-references academic integrity standards and clarifies how AI-enabled misconduct is interpreted within existing unauthorized assistance language, due process expectations, and evidentiary practices. This is especially important given widespread caution about the limits of AI detection tools in misconduct decisions (Cotton et al., 2024).

Coherence can be strengthened without rewriting entire integrity codes by aligning language across student handbooks, integrity procedures, IT guidance, and teaching center resources, and by providing consistent examples of permissible and impermissible use that can be adapted locally.

Treat Equity as a Design Requirement

When acceptable use is determined primarily through classroom policy, variability is inevitable. Equity concerns arise when variability becomes unpredictability or when students must navigate ambiguous expectations without consistent supports. Institutional leaders should reduce equity risk by standardizing how expectations are communicated to students across courses, such as adopting shared templates for syllabi and LMS pages, defining what “ask first” means in practice, and providing plain-language disclosure models that students can follow. Leaders should also examine access implications of recommended tools, including whether tools require paid accounts, whether campus-provided alternatives exist, and whether privacy restrictions create uneven access to beneficial supports. Equity-forward governance in a decentralized model depends on clear communication and accessible infrastructure, not only statements of principle.

Invest in Institutional Capacity for Implementation

Because instructors are often the primary determiners of classroom AI rules, faculty support should be treated as a governance strategy, not an optional add-on. In practice, this means building structures that better coordinate academic authority with administrative and compliance-oriented functions. Leaders should invest in professional learning that is discipline-aware and assessment-focused, including applied exemplars of assignment redesign, documentation practices, and learning-centered uses of AI. Teaching and learning centers can be resourced to provide ready-to-use policy language and decision supports that reduce the time burden on individual instructors. Institutions should also strengthen cross-unit coordination among academic affairs, IT, legal, libraries, disability services, and student support units so that instructional guidance aligns with tool vetting, privacy expectations, procurement constraints, and accommodations.

Institutionalize Review and Feedback Cycles

Generative AI governance will remain iterative. Institutions should adopt predictable review cycles for approved tools guidance, data classification

messaging, sample syllabus language, and student-facing resources. Leaders can strengthen legitimacy and usability by building feedback loops that capture faculty and student experiences with policy ambiguity, enforcement challenges, and support needs. A defined cadence for review also reduces the risk that guidance fragments across disconnected webpages and units. Over time, these mechanisms can support gradual policy maturation while preserving the flexibility that institutions currently value.

Conclusions

This study provides a comparative policy mapping of publicly available generative AI governance artifacts across America's 50 flagship public universities. The findings depict a consistent institutional approach in which instructional governance is largely distributed. Universities commonly rely on guidance and resources while delegating classroom rule-setting to instructors, and they concentrate institution-level authority in the domains of data protection, approved tools, and risk management. This governance architecture helps institutions preserve pedagogical flexibility and respond to rapid technological change, but it also produces predictable implementation challenges when expectations vary widely across courses and programs.

For instructors, the findings highlight a dual reality of autonomy and burden. Faculty are empowered to determine classroom AI rules and to align those rules with course learning goals. At the same time, they are expected to translate evolving institutional guidance into enforceable classroom policy, communicate expectations clearly, and design assessments that support integrity in an environment where detection tools are widely treated with caution (McDonald et al., 2025). For students, the distributed model frames responsible use through norms of clarification and transparency, including expectations to confirm course rules and disclose permitted AI use. However, when course policies vary substantially, students may experience inconsistent expectations and uneven enforcement, which can undermine clarity and trust and raise equity concerns in how policies are understood and applied.

More broadly, the analysis suggests that current governance efforts are still in a maturation phase. Institutions have built substantial guardrails around data privacy and tool governance, but instructional policy remains largely localized and dependent on faculty capacity, support structures, and alignment across campus units. Sector-level evidence also suggests that policy infrastructure is developing unevenly, reinforcing the need for leadership attention to coherence, capacity, and student-facing clarity as AI use expands (Robert & McCormack, 2025). Viewed through Clark's (1983) Triangle of Coordination, this convergence reflects a relatively stable equilibrium across

academic, state-administrative, and market domains, suggesting that major departures from the current governance pattern are more likely to emerge from broader external pressures than from internal policy redesign alone. This study contributes to educational policy scholarship by offering a baseline map of how flagship public universities are currently structuring generative AI governance and by identifying recurring governance choices that shape implementation at the classroom level. The coding indicators and thematic patterns can support future benchmarking and comparative work, including research that examines how governance approaches evolve over time and how institutional postures influence faculty practice, student outcomes, and equity.

Future research should move beyond document analysis to examine policy enactment and implementation in practice. Qualitative studies of how faculty, students, and administrators interpret and apply guidance would help clarify how distributed governance operates on the ground and which supports reduce inconsistency and inequity. Longitudinal analyses could track how policies shift as tools change, legal expectations develop, and institutions learn from early implementation. Comparative research across institutional types, including community colleges, private institutions, and international contexts, would further strengthen understanding of how governance structures shape AI policy choices and their consequences.

ORCID iD

Jason LaFrance  <https://orcid.org/0000-0002-0636-4193>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

- Airaj, M. (2024). Ethical artificial intelligence for teaching-learning in higher education. *Education and Information Technologies*, 29, 17145–17167. <https://doi.org/10.1007/s10639-024-12545-x>
- Al-Zahrani, A. M., & Alasmari, T. M. (2024). Exploring the impact of artificial intelligence on higher education: The dynamics of ethical, social, and educational implications. *Humanities and Social Sciences Communications*, 11(1), 1–12. <https://doi.org/10.1057/s41599-024-03432-4>

- An, Y., Yu, J., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22, Article 10. <https://doi.org/10.1186/s41239-025-00507-3>
- Barany, A., Nasiar, N., Porter, C., Zambrano, A. F., Andres, A. L., Bright, D., Shah, M., Liu, X., Gao, S., Zhang, J., Mehta, S., Choi, J., Giordano, C., Baker, R. S., & Baker, R. S. (2024, July). ChatGPT for education research: Exploring the potential of large language models for qualitative codebook development. In O. C. Santos, H. Rodrigo, & I. I. Bittencourt (Eds.), *International conference on artificial intelligence in education* (pp. 134–149). Springer Nature Switzerland.
- Batista, J., Mesquita, A., & Carnaz, G. (2024). Generative AI and higher education: Trends, challenges, and future directions from a systematic literature review. *Information*, 15(11), Article 676. <https://doi.org/10.3390/info15110676>
- Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27–40. <https://doi.org/10.3316/QRJ0902027>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- 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), Article 38. <https://doi.org/10.1186/s41239-023-00408-3>
- Clark, B. R. (1983). *The higher education system: Academic organization in cross-national perspective*. University of California Press.
- 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>
- Creswell, J. W. (2021). *Educational research: Planning, conducting, and evaluating quantitative and qualitative research* (6th ed.). Pearson.
- Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). Sage Publications.
- Dabis, A., & Csáki, C. (2024). AI and ethics: Investigating the first policy responses of higher education institutions to the challenge of generative AI. *Humanities and Social Sciences Communications*, 11(1), 1–13. <https://doi.org/10.1057/s41599-024-03526-z>
- Dignum, V. (2020). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer.
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160. <https://doi.org/10.4324/9781315247533-34>
- Dotan, R., Parker, L. S., & Radzilowicz, J. (2024, June). Responsible adoption of generative AI in higher education: Developing a “points to consider” approach based on faculty perspectives. In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency* (pp. 2033–2046). <https://doi.org/10.1145/3630106.3659023>

- Ganguly, A., Johri, A., Ali, A., & McDonald, N. (2025). Generative artificial intelligence for academic research: Evidence from guidance issued for researchers by higher education institutions in the United States. *AI and Ethics*, 5(4), 3917–3933. <https://doi.org/10.1007/s43681-025-00688-7>
- Ghimire, A., & Edwards, J. (2024, July). From guidelines to governance: A study of AI policies in education. In *International conference on artificial intelligence in education* (pp. 299–307). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-64312-5_36
- Gianni, R., Lehtinen, S., & Nieminen, M. (2022). Governance of responsible AI: From ethical guidelines to cooperative policies. *Frontiers in Computer Science*, 4, Article 873437. <https://doi.org/10.3389/fcomp.2022.873437>
- International Association of Privacy Professionals (IAPP) & FTI Consulting. (2024). *AI governance in practice report 2024*. <https://iapp.org/resources/article/ai-governance-in-practice-report/>
- Kingdon, J. W. (1995). *Agendas, alternatives, and public policies* (2nd ed.). HarperCollins.
- LaFrance, J. (2025a). The power and perils of AI-driven assessment in higher education. In M. Kessinger, G. Nelson, L. Lennex, & K. Nettleton (Eds.), *Examining AI disruption in educational settings: Challenges and opportunities* (pp. 267–296). IGI Global Scientific Publishing. <https://doi.org/10.4018/979-8-3693-7873-1.ch010>
- LaFrance, J. (2025b). Navigating educational leadership research with generative artificial intelligence: Innovations, ethical considerations, and applications. In J. Bailey Watters, G. Miller, & R. Rhone (Eds.), *Reimagining the P-20 landscape for school leadership learning* (pp. 139–168). IGI Global Scientific Publishing. <https://doi.org/10.4018/979-8-3693-6220-4.ch006>
- Longo, L. (2019, September). Empowering qualitative research methods in education with artificial intelligence. In *World conference on qualitative research* (pp. 1–21). Springer International Publishing.
- Luckin, R. (2018). *Machine learning and human intelligence: The future of education for the 21st century*. UCL Institute of Education Press.
- Marshall, D. T., & Naff, D. B. (2024). The ethics of using artificial intelligence in qualitative research. *Journal of Empirical Research on Human Research Ethics*, 19(3), 92–102. <https://doi.org/10.31235/osf.io/3rnbh>
- McDonald, N., Johri, A., Ali, A., & Hingle-Collier, A. (2025). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. *Computers in Human Behavior: Artificial Humans*, 3, Article 100121. <https://doi.org/10.1016/j.chbah.2025.100121>
- Merriam, S. B. (2009). *Qualitative research: A guide to design and implementation* (2nd ed.). Jossey-Bass.
- Michel-Villarreal, R., Vilalta-Perdomo, E., Salinas-Navarro, D. E., Thierry-Aguilera, R., & Gerardou, F. S. (2023). Challenges and opportunities of generative AI for higher education as explained by ChatGPT. *Education Sciences*, 13(9), Article 856. <https://doi.org/10.3390/educsci13090856>

- Morgan, D. L. (2023). Exploring the use of artificial intelligence for qualitative data analysis: The case of ChatGPT. *International Journal of Qualitative Methods*, 22, 16094069231211248. <https://doi.org/10.1177/16094069231211248>
- National Center for Education Statistics. (2023). *Locale boundaries and definitions* (NCES 2023-125). U.S. Department of Education, Institute of Education Sciences. https://nces.ed.gov/programs/edge/docs/EDGE_NCES_LOCALE.pdf
- National Center for Education Statistics. (2024). *College Navigator*. U.S. Department of Education, Institute of Education Sciences. <https://nces.ed.gov/collegenavigator/>
- Patton, M. Q. (2015). *Qualitative research & evaluation methods: Integrating theory and practice* (4th ed.). Sage Publications.
- Popenici, S. A., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12(1), Article 22. <https://doi.org/10.1186/s41039-017-0062-8>
- Ren, K., & Li, J. (2013). Academic freedom and university autonomy: A higher education policy perspective. *Higher Education Policy*, 26, 507–522. <https://doi.org/10.1057/hep.2013.31>
- Robert, J. (2024, February 12). *2024 EDUCAUSE AI Landscape Study*. EDUCAUSE. <https://library.educause.edu/resources/2024/2/2024-educause-ai-landscape-study>
- Robert, J., & McCormack, M. (2025, February 17). *2025 EDUCAUSE AI Landscape Study*. EDUCAUSE. <https://library.educause.edu/resources/2025/2/2025-educause-ai-landscape-study>
- Rizvi, F., & Lingard, B. (2010). *Globalizing education policy*. Routledge.
- Saldaña, J. (2021). *The coding manual for qualitative researchers* (4th ed.). Sage Publications.
- Slimi, Z., & Carballido, B. V. (2023). Navigating the ethical challenges of artificial intelligence in higher education: An analysis of seven global AI ethics policies. *TEM Journal: Technology, Education, Management, Informatics*, 12(2), 590–602. <https://doi.org/10.18421/tem122-02>
- Southworth, J., Migliaccio, K., Glover, J., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for AI across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, 4, Article 100127. <https://doi.org/10.1016/j.caeai.2023.100127>
- Spivakovsky, O. V., Omelchuk, S. A., Kobets, V. V., Valko, N. V., & Malchykova, D. S. (2023). Institutional policies on artificial intelligence in university learning, teaching and research. *Information Technologies and Learning Tools*, 97(5), 181–202. <https://doi.org/10.33407/itlt.v97i5.5395>
- Theelen, H., Vreuls, J., & Rutten, J. (2024). Doing research with help from ChatGPT: Promising examples for coding and inter-rater reliability. *International Journal of Technology in Education*, 7(1), 1–18. <https://doi.org/10.46328/ijte.537>

- Voigt, P., & von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR): A practical guide*. Springer. <https://doi.org/10.1007/978-3-319-57959-7>
- Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in higher education: Seeing ChatGPT through universities' policies, resources, and guidelines. *Computers and Education: Artificial Intelligence*, 7, Article 100326. <https://doi.org/10.1016/j.caeai.2024.100326>
- Xiao, P., Chen, Y., & Bao, W. (2023). *Waiting, banning, and embracing: An empirical analysis of adapting policies for generative AI in higher education* [SSRN working paper]. SSRN. <https://doi.org/10.2139/ssrn.4458269>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). Sage Publications.

Author Biography

Jason LaFrance is an Associate Professor of Educational Leadership at Appalachian State University. His research focuses on educational governance, academic integrity, and technology leadership in higher education, with particular attention to artificial intelligence policy and institutional decision-making. His applied scholarship examines how AI is enacted in practice, including its use in instructional design, faculty work, and student learning through the Scholarship of Teaching and Learning (SoTL). Drawing on prior K-12 leadership experience, Dr. LaFrance teaches master's and doctoral levels in educational leadership and organizational improvement.

Appendix A

Table A1. Flagship Universities in the United States.

State	Flagship university	Campus setting	Undergraduate enrollment	Graduate enrollment	Carnegie classification (R1/R2)
Alabama	University of Alabama	City: Midsize	33,435	6,187	Doctoral Universities: Very High Research Activity (R1)
Alaska	University of Alaska Fairbanks	City: Small	5,638	1,002	Doctoral Universities: High Research Activity (R2)
Arizona	University of Arizona	City: Large	41,899	11,102	R1
Arkansas	University of Arkansas	City: Small	27,472	4,668	R1
California	University of California, Berkeley	City: Midsize	33,078	12,621	R1
Colorado	University of Colorado Boulder	City: Midsize	32,105	9,327	R1
Connecticut	University of Connecticut	Town: Fringe	19,388	7,976	R1
Delaware	University of Delaware	Suburb: Large	19,772	4,449	R1
Florida	University of Florida	City: Midsize	34,924	19,890	R1
Georgia	University of Georgia	City: Midsize	31,514	10,101	R1
Hawaii	University of Hawaii at Manoa	City: Large	14,576	4,680	R1
Idaho	University of Idaho	Town: Distant	9,269	2,580	R2
Illinois	University of Illinois, Urbana-Champaign	City: Small	35,564	20,999	R1
Indiana	Indiana University Bloomington	City: Small	36,833	10,694	R1
Iowa	University of Iowa	City: Small	22,130	7,912	R1
Kansas	University of Kansas	City: Small	28,406	7,710	R1
Kentucky	University of Kentucky	City: Large	23,930	8,773	R1
Louisiana	Louisiana State University, Baton Rouge	City: Midsize	32,574	6,844	R1
Maine	University of Maine	Suburb: Small	9,267	2,493	R1
Maryland	University of Maryland, College Park	Suburb: Large	30,608	10,205	R1
Massachusetts	University of Massachusetts Amherst	City: Small	23,936	7,874	R1
Michigan	University of Michigan, Ann Arbor	City: Midsize	33,730	18,335	R1
Minnesota	University of Minnesota Twin Cities	City: Large	39,558	15,334	R1
Mississippi	University of Mississippi	Town: Remote	19,094	4,949	R1

(continued)

Table A1. (continued)

State	Flagship university	Campus setting	Undergraduate enrollment	Graduate enrollment	Carnegie classification (R1/R2)
Missouri	University of Missouri-Columbia	City: Midsize	23,613	7,400	R1
Montana	University of Montana	City: Small	7,570	2,757	R1
Nebraska	University of Nebraska, Lincoln	City: Large	19,338	4,648	R1
Nevada	University of Nevada, Reno	City: Large	17,965	3,813	R1
New Hampshire	University of New Hampshire	Town: Fringe	11,376	2,250	R1
New Jersey	Rutgers University, New Brunswick	City: Small	36,588	14,029	R1
New Mexico	University of New Mexico	City: Large	16,905	5,795	R1
New York	University at Buffalo	Suburb: Large	20,463	11,426	R1
North Carolina	University of North Carolina at Chapel Hill	City: Small	20,681	11,553	R1
North Dakota	University of North Dakota	City: Midsize	9,791	2,161	R1
Ohio	Ohio State University	City: Large	45,728	14,318	R1
Oklahoma	University of Oklahoma, Norman	Suburb: Midsize	22,025	7,120	R1
Oregon	University of Oregon	City: Midsize	19,963	3,823	R1
Pennsylvania	Pennsylvania State University	City: Small	42,223	8,176	R1
Rhode Island	University of Rhode Island	Suburb: Large	14,670	2,681	R2
South Carolina	University of South Carolina, Columbia	City: Midsize	28,470	8,109	R1
South Dakota	University of South Dakota	Town: Distant	7,095	2,773	R2
Tennessee	University of Tennessee, Knoxville	City: Midsize	28,883	7,421	R1
Texas	University of Texas at Austin	City: Large	42,444	10,638	R1
Utah	University of Utah	City: Midsize	26,827	8,433	R1
Vermont	University of Vermont	City: Small	12,168	2,152	R2
Virginia	University of Virginia	Suburb: Midsize	17,612	8,312	R1
Washington	University of Washington, Seattle	City: Large	39,515	16,105	R1
West Virginia	West Virginia University	City: Small	18,615	5,585	R1
Wisconsin	University of Wisconsin, Madison	City: Large	36,797	12,808	R1
Wyoming	University of Wyoming	Town: Remote	8,250	2,663	R2

Note. Campus settings terminology defined in Table B1.

Source. Data from National Center for Education Statistics (NCES, 2024) College Navigator (campus setting and enrollments as of Fall 2023) and Carnegie Classification (R1/R2).

Appendix B

Table B1. Campus Setting Definitions.

Setting	Definition
City: Large	Territory inside an urbanized area and inside a principal city with a population of 250,000 or more.
City: Midsize	Territory inside an urbanized area and inside a principal city with a population between 100,000 and 250,000.
City: Small	Territory inside an urbanized area and inside a principal city with a population less than 100,000.
Suburb: Large	Territory outside a principal city and inside an urbanized area with a population of 250,000 or more.
Suburb: Midsize	Territory outside a principal city and inside an urbanized area with a population between 100,000 and 250,000.
Suburb: Small	Territory outside a principal city and inside an urbanized area with a population less than 100,000.
Town: Fringe	Territory inside an urban cluster that is less than or equal to 10 miles from an urbanized area.
Town: Distant	Territory inside an urban cluster that is more than 10 but less than or equal to 35 miles from an urbanized area.
Town: Remote	Territory inside an urban cluster that is more than 35 miles from an urbanized area.
Rural: Fringe	Census-defined rural territory that is less than or equal to 5 miles from an urbanized area, or less than or equal to 2.5 miles from an urban cluster.
Rural: Distant	Rural territory that is more than 5 but less than or equal to 25 miles from an urbanized area, or more than 2.5 but less than or equal to 10 miles from an urban cluster.
Rural: Remote	Rural territory that is more than 25 miles from an urbanized area and more than 10 miles from an urban cluster.

Source. Setting and definition established by National Center for Education Statistics (NCES, 2023).

Appendix C

Table C1. Coding Indicators and Descriptive Frequencies.

Policy feature	Coding criterion	n (of 50)	%
Course-level discretion	Course policy determined at instructor/course level (e.g., syllabus guidance expected)	50	100
Sample syllabus statements provided	Institution provides sample syllabus language/templates	44	88
Explicit AI language in academic integrity codes	Academic integrity policy explicitly references AI/GAI	17	34
Discrete institutional AI guidance from technology department	Separate AI guidance page/memo (e.g., IT guidance distinct from integrity code)	42	84
Data/tool governance guidance	Approved tools list and/or data classification guidance for AI use	46	92
Mandatory use of approved tools	Institution requires use of vetted/approved tools for certain contexts/data	7	14
Research integrity/compliance mentioned	Research guidance includes compliance/integrity expectations for AI use	33	66
Disclosure required	Policy/guidance mandates disclosure of AI use in coursework	15	30
Disclosure recommended	Policy/guidance recommends disclosure (but does not mandate)	34	68
Equity/access mentioned	Policy/guidance references equity, access, inclusion, or accommodations	19	38
Bias/fairness mentioned	Policy/guidance references bias, fairness, discrimination, or related risks of AI	49	98
Policy posture: Balanced	Coded dominant posture = guidance-oriented/conditional	32	64
Policy posture: Transformative	Coded dominant posture = innovation/AI literacy/ethics-forward	10	20
Policy posture: Restrictive	Coded dominant posture = prohibitive/compliance-forward	8	16

Note: Indicators were coded from primary institutional governance artifacts; counts reflect institutions (n=50), not total documents.

Appendix D

Table D1. Stakeholder-Facing Guidance Examples.

Audience	Guidance element	Description	Representative examples	Implementation implications
Faculty	Syllabus statements	Institutions provide sample clauses with varying stances (prohibited, conditional, permissive).	Arkansas: prohibited/restricted/unrestricted options; Michigan: encourage/allow with disclosure/prohibit.	Supports consistent course communication; reduces ambiguity for students.
Faculty	Pedagogical prompts	Guidance encourages alignment of AI use with learning outcomes and disciplinary goals.	Arizona: alignment with learning outcomes; Kansas: include rationale for policy choice.	Promotes intentional policy design rather than reactive rules.
Faculty	Teaching center resources	Teaching and learning centers provide frameworks, workshops, and templates.	Georgia: learn-explore-reflect process before setting stance.	Builds faculty capacity and supports implementation.
Faculty	Assessment design	Institutions encourage assessment redesign to foreground process and higher-order thinking.	Michigan: protect “cognitive dimension of learning” and student responsibility.	Reduces reliance on detection; strengthens integrity through design.
Students	“Ask first”	Students are advised to confirm expectations because policies vary by course/instructor.	Colorado Boulder: ask instructors; Washington: permission-based expectations.	Encourages clarification; can create uneven experiences across courses if expectations vary widely.
Students	“Disclose”	Students are expected to document permitted AI use (required or recommended depending on institution/course).	Arkansas: attribution in permissive model; Michigan: responsibility for facts/citations.	Reinforces transparency and academic honesty norms.
Students	Privacy reminders	Guidance discourages inputting sensitive/protected data into public AI tools.	Florida: protect personal/protected info; Minnesota: restrictions for private/restricted data.	Builds awareness of privacy/FERPA-related responsibilities.
Students	Framing of AI use	AI is framed as a tool to support learning rather than replace it.	UNC Chapel Hill: “AI should help you think. Not think for you.”	Reinforces human accountability and learning-centered use.

Note: Table D1 provides illustrative examples of common guidance elements and is not intended to be exhaustive.