Educational and Humanities
ISSN Print: 3104-4352
ISSN Online: 3104-4360
CD ISSN: 3104-4379
https://doi.org/10.63313/EH.9055

Generative AI in Civic Education: A Multi-Case Study of Pedagogical Innovation and Governance in Higher Education

ShengYu Gu
School of Geography and Tourism, Huizhou University, Huizhou, China
Email: miller@hzu.edu.cn

How to cite this paper: Gu, S. Y. (2026). Generative AI in civic education: A multi-case study of pedagogical innovation and governance in higher education. *Educational and Humanities*, 3(1), 9–30. ISSN Print: 3104-4352; ISSN Online: 3104-4360.
https://doi.org/10.63313/EH.9055
Published: 2026-04-16

Copyright © 2026 by author(s) and Erytis Publishing Limited.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/
CC BY 4.0

Abstract

The rapid diffusion of generative artificial intelligence (AI) is reshaping higher education, raising critical questions about its integration into civic reasoning, public judgment, and ethical learning. While existing studies have examined its pedagogical opportunities and risks, limited attention has been paid to practice-oriented institutional models that are both educationally effective and transferable. This study addresses this gap by examining how universities operationalize generative AI through the interaction of technology, pedagogy, and governance, and by identifying lessons for emerging institutions. Using a qualitative multiple-case study design, the study analyzes three international cases: Harvard University (USA), University College London (UK), and the University of Helsinki (Finland). Data were collected from institutional reports, policy documents, academic publications, and open-access course materials, and analyzed through thematic analysis. Drawing on digital governance theory, experiential learning theory, and political socialization in digital environments, the study develops a Technology-Pedagogy-Governance (TPG) framework for comparative analysis. The findings reveal three distinct but complementary institutional pathways. Harvard demonstrates AI-supported civic reasoning through its use as an analytical assistant; UCL highlights governance-centered design through differentiated assessment frameworks; and Helsinki emphasizes AI literacy and ethics as a foundation for civic competence through scalable open-access education. Across cases, effective practices converge on controlled AI integration, critical thinking-centered pedagogy, and explicit ethical governance.
This study contributes by extending generative AI research into civic education, advancing institution-based analysis beyond abstract debates, and proposing a transferable AI-enhanced civic education model. It suggests that emerging universities should adopt a sequenced approach: prioritizing AI literacy, followed by bounded pedagogical integration, and supported by clear governance frameworks.

Keywords

Generative artificial intelligence; Higher education; Civic education; AI governance

1. Introduction

1.1. Background: Generative AI and the Transformation of Higher Education

The public release of large language models such as ChatGPT has accelerated a structural shift in higher education. Since 2023, artificial intelligence in education has entered a new phase in which generative systems not only automate routine tasks but also produce text, arguments, code, and feedback directly applicable to teaching and learning. This rapid development has reshaped pedagogy, prompting universities to reconsider assessment, academic integrity, curriculum design, and digital literacy (Safdar et al., 2026; Tlili, 2025; Mochizuki et al., 2025). This transformation builds on earlier AI expansion in higher education. A review of 138 studies found a sharp increase in AI-related research in 2021-2022, covering assessment, tutoring, prediction, and learning management (Crompton & Burke, 2023). Generative AI thus intensifies existing debates while introducing new concerns regarding authorship, cognitive offloading, and epistemic trust.

Institutional responses have been equally rapid. UNESCO emphasizes the need for immediate policy action and human-centred governance (Mochizuki et al., 2025), while the OECD reports that multiple jurisdictions have introduced soft governance measures, though dedicated regulation remains limited (Bo, 2025). These developments indicate that generative AI presents a systemic challenge extending beyond classroom experimentation to institutional governance. The implications are particularly significant for civic and political education. Digital platforms already shape political learning, and increased social media use has been linked to higher political knowledge (Kusumawardani et al., 2025). Generative AI further transforms this environment by enabling students to co-produce arguments and interpretations through systems that appear authoritative but may lack reliability or transparency. Universities are also emerging as sites of institutional experimentation. Arizona State University has pursued enterprise-level AI integration (Shaikh et al., 2026), University College London has developed structured governance frameworks for assessment (Gupta & Nyamapfene, 2025), and the University of Helsinki has expanded global AI literacy through its Elements of AI initiative (Bezeriță, 2025). These diverse responses highlight distinct institutional pathways, forming the basis for the comparative analysis developed in later chapters.

1.2. Problem Statement: Gaps in AI-Driven Civic and Political Education

Despite rapid growth in the literature, a key gap remains. Existing research focuses primarily on general teaching applications such as writing support, feedback, and assessment (Safdar et al., 2026; Tlili, 2025), while paying limited attention to civic and political education. The central issue is not only whether generative AI improves efficiency, but whether it can support civic reasoning, ethical reflection, and informed participation without undermining these capacities. This gap is

significant for three reasons. First, civic education depends on interpretive judgment, requiring students to evaluate evidence, compare arguments, and engage with uncertainty. However, AI-generated outputs may appear coherent regardless of accuracy, making critical oversight essential (Safdar et al., 2026). Second, governance frameworks remain uneven. OECD findings show that policy responses are still fragmented (Bo, 2025). Institutional models such as UCL's differentiated assessment framework demonstrate promising approaches, yet evidence on their effectiveness in civic-oriented learning remains limited.

Third, transferability is underexplored. Much existing research focuses on well-resourced institutions, offering limited guidance for universities operating under constrained conditions. The key issue is not replication, but adaptation-identifying which elements of international practice are feasible across diverse contexts. This study addresses this gap through a comparative case approach.

1.3. Research Aim: Learning Transferable Models from International Cases

This study aims to identify transferable institutional models for integrating generative AI into civic and political education. Rather than proposing a universal framework, it examines how universities configure technology, pedagogy, and governance, and which elements can be adapted by emerging institutions. A case-study design is appropriate because generative AI is embedded in institutional contexts (Helström, 2026). The study focuses on three contrasting cases: enterprise-level integration (Arizona State University), governance-centered implementation (University College London), and large-scale AI literacy (University of Helsinki). These cases enable comparative analysis of distinct institutional pathways. The contribution is both theoretical and practical. Theoretically, the study connects AI in education with civic learning and digital governance. Practically, it develops a transferable analytical model-elaborated in Chapter 3 and synthesized in Chapter 6-to guide institutions in adopting AI in ways that are pedagogically effective and ethically grounded.

1.4. Research Questions

This study addresses three research questions:

RQ1. How is generative AI used in civic and political education internationally ?

Examines institutional forms of AI adoption, including classroom use, AI literacy, and learning design.

RQ2. What governance and pedagogical mechanisms ensure effectiveness ?

Focuses on assessment rules, academic integrity, disclosure, and critical literacy.

RQ3. What lessons are transferable to emerging universities ?

Identifies adaptable practices across different institutional contexts.

These questions guide the empirical analysis in Chapter 5 and the model development in Chapter 6.

1.5. Structure of the Paper

The paper is organized as follows. Chapter 2 reviews relevant literature and identifies research gaps. Chapter 3 presents the theoretical and analytical framework. Chapter 4 outlines the methodology and case selection. Chapter 5 provides case analysis and comparison. Chapter 6 synthesizes findings and develops a transferable model for institutional application.

2. Literature Review

2.1. Generative AI in Higher Education: Pedagogical Applications

Research on generative artificial intelligence (AI) in higher education has expanded rapidly, building on earlier work on AI in education. A review of 138 studies shows that research increased sharply in 2021-2022, focusing on assessment, tutoring, and learning management (Crompton & Burke, 2023). Generative AI therefore represents an acceleration rather than a rupture, distinguished by the immediacy and accessibility of AI-generated outputs in academic work. Safdar et al. (2026) provide a key synthesis, identifying applications such as writing support, feedback, and tutoring, while emphasizing risks including hallucination, bias, and overreliance. Subsequent studies reinforce this dual perspective. Tlili (2025) highlights benefits for learning support but raises concerns about dependency and assessment validity, while Cotton et al. (2024) argue that AI-assisted writing may reshape how critical thinking is developed rather than merely facilitating misconduct. A further strand critiques underlying pedagogical assumptions. Liu and Glave (2026) argue that AI-driven "personalization" often reflects standardized learning pathways rather than genuine responsiveness. This is particularly relevant for civic and political education, where pedagogical value depends on fostering inquiry, pluralism, and reflexive judgment rather than efficiency alone. Overall, two limitations remain. First, the literature rarely distinguishes between general AI use and applications in civic-oriented education. Second, it remains largely tool-centred, offering limited insight into how institutions organize pedagogy, assessment, and accountability. These gaps justify the present study's focus on institutional case analysis (see Sections 2.4 and Chapters 4-5).

2.2. Civic and Political Education in Digital Contexts

The literature on civic and political education in digital environments provides a foundation for understanding the implications of generative AI. Bowyer and Kahne (2020) show that different digital learning opportunities produce distinct civic outcomes: content creation supports online participation, while content evaluation supports offline engagement. This suggests that design, rather than mere access, determines educational impact. Similarly, Kusumawardani et al. (2021) find that

social media use enhances political knowledge through online discussion, indicating that digital interaction plays a constitutive role in political learning. Generative AI extends this environment by enabling students not only to access and discuss information, but also to co-produce arguments and interpretations. Recent research further highlights conceptual shifts in civic education. Alam (2026) shows that citizenship learning is increasingly distributed across formal and informal environments, including social media and peer networks. As a result, civic education is no longer confined to the classroom but embedded in hybrid, platform-mediated ecosystems. However, existing scholarship has not sufficiently addressed how generative AI should be integrated into this context. While digital civic education research emphasizes media literacy and participation, it has yet to fully account for AI systems that generate persuasive, personalized, and potentially misleading content. This study extends that literature by examining AI as a new mediator of civic learning.

2.3. AI Ethics, Governance, and Algorithmic Risks

A third body of literature focuses on ethics and governance. Here, there is broad consensus that AI adoption must address issues of fairness, transparency, privacy, and accountability. UNESCO's guidance frames generative AI as a governance challenge, emphasizing the need for human-centred approaches and institutional oversight (Mochizuki et al., 2025). Nguyen (2025) identifies recurring ethical principles, including transparency, justice, and human agency, demonstrating that ethical concerns have been institutionalized in policy discourse. However, implementation remains uneven, as institutions often lack concrete procedures for integrating these principles into practice. Holmes and colleagues further argue that AI systems embed implicit assumptions about learning, warning against equating personalization with meaningful educational development. These concerns are particularly significant for civic education, where algorithmic bias may shape not only learning outcomes but also public reasoning and knowledge legitimacy. Governance research also highlights institutional variation. OECD findings indicate that policy responses remain fragmented (Bo, 2025), suggesting that governance must be examined at the institutional level. In practice, students encounter AI through local rules governing assessment, disclosure, and acceptable use. This reinforces the need for institution-based analysis, motivating the case-study approach adopted in this study.

2.4. Case Study Approach in Educational Technology Research

Case study research is well suited to examining generative AI in higher education because the phenomenon is inseparable from its institutional context. Helström (2026) emphasizes that case studies are appropriate for contemporary, context-dependent phenomena, a condition that clearly applies to AI in education. In educational research, case studies enable analysis of implementation and

organizational processes rather than isolated variables. Merriam's work highlights their value for in-depth investigation of bounded settings, while Lee (2025) and Mott and Haddad (2025) show that different traditions converge on the importance of contextual understanding. This is particularly relevant for generative AI, whose educational effects vary depending on institutional arrangements. A tool may function as a tutor, writing assistant, or governance challenge depending on context. Multi-case comparison therefore allows analysis of how technology, pedagogy, and governance interact, rather than simply whether AI is used. Accordingly, this study adopts a comparative case-study design to identify transferable institutional logics, as developed in Chapter 6.

Table 1. Key International Studies on AI in Civic and Political Education

Author/Year	Context	AI Application	Key Findings
Safdar et al. (2026)	Europe / higher education debate	ChatGPT and large language models in education	Identifies substantial pedagogical potential, including tutoring and feedback, but warns that hallucinations, bias, and overreliance require critical oversight.
Cotton et al. (2024)	United Kingdom / higher education	AI-assisted writing and academic work	Shows that generative AI challenges academic integrity and may alter how critical thinking and original writing are developed and assessed.
Mochizuki et al. (2025)	Global policy context	Generative AI in education and research	Advocates a human-centred governance framework emphasizing privacy protection, ethical validation, regulation, and pedagogical design.
Liu and Glave (2026)	Global / AI in education	AI tutoring, adaptive systems, and educational design	Argues that AI may support personalization, but many systems rest on narrow pedagogical assumptions and require critical governance.

2.5. Research Gap: Lack of Practice-Oriented, Transferable Institutional Models

The literature is substantial but incomplete. Pedagogical studies demonstrate AI's potential but remain general; civic education research highlights digital transformation but largely predates generative AI; and governance studies articulate principles without fully operationalizing them at the institutional level. Three gaps follow. First, research on generative AI in civic and political education remains limited. Second, existing studies emphasize opportunities and risks rather than institutional arrangements. Third, there is a lack of transferable models for universities operating under constrained conditions. This study addresses these gaps through a multi-case design focused on institutional practice. By examining how universities integrate technology, pedagogy, and governance, it moves beyond abstract debates toward actionable, transferable insights, as developed in Chapters 3-6.

3. Theoretical and Analytical Framework

3.1. Theoretical Foundations

This study draws on three complementary perspectives: digital governance theory, experiential learning theory, and political socialization in digital environments. Together, they explain why generative AI in higher education should not be

understood merely as a technical innovation. Its significance depends on how universities govern its use, how students encounter it in learning processes, and how digital interaction shapes civic understanding. Building on Chapter 2, this chapter develops the conceptual basis for the comparative case analysis in Chapter 5.

3.1.1. Digital Governance Theory

Digital governance theory provides the first pillar because generative AI is not only a pedagogical tool but part of a broader transformation in how institutions organize information and regulate socio-technical systems. Digital technologies reshape public institutions through reintegration, needs-based organization, and digitization, displacing older bureaucratic routines. Although developed outside higher education, this perspective is directly relevant to universities, where AI adoption requires changes not only in classroom practice but also in assessment, staff guidance, student support, data protection, and oversight. Universities therefore do not simply use AI; they govern its conditions of use. This theory is useful in two ways. First, it shifts attention from technological features to the institutional arrangements that make technology legitimate and functional, consistent with the governance concerns discussed in Section 2.3. Second, it highlights that governance is not only external regulation but also internal policy architecture through which institutions define acceptable conduct, allocate responsibility, and align digital tools with educational aims. In this study, digital governance theory frames universities as active governors of AI-mediated learning. It helps explain why the selected cases differ not only in technology but also in governance logic: Arizona State University emphasizes institutional experimentation, University College London emphasizes formalized assessment governance, and the University of Helsinki emphasizes public-facing AI literacy. These are distinct institutional responses to the same technological moment.

3.1.2. Experiential Learning Theory

The second pillar is experiential learning theory, associated with Kolb's model of learning as a cycle of concrete experience, reflective observation, abstract conceptualization, and active experimentation. This perspective is useful because it treats learning as a dynamic process of engagement, reflection, and revision rather than passive reception. Experiential learning theory is relevant to generative AI for two reasons. First, generative AI enters directly into the learning cycle by shaping how students draft arguments, test claims, generate examples, or receive feedback. It may support inquiry, but it may also compress reflection. A student using AI to refine a policy argument may benefit from rapid iteration while losing the reflective struggle necessary for deep learning. This aligns with concerns raised in Section 2.1 by Safdar et al. (2026) and Cotton et al. (2024), both of whom argue that AI is educationally valuable only when accompanied by critical engagement. Second, the

theory is highly compatible with civic and political education because democratic learning is inherently experiential. Students develop civic judgment through debate, evidence evaluation, and engagement with disagreement. Generative AI may enhance these processes through simulation and iterative questioning, but it may also weaken them if students rely uncritically on polished outputs. The key issue, therefore, is whether pedagogical design preserves reflection and inquiry rather than replacing them with convenience.

3.1.3. Political Socialization in Digital Environments

The third pillar is political socialization in digital environments. Classical political socialization theory emphasizes that political orientations are formed through interaction with institutions, media, and peers. In contemporary settings, digital media have become central to that process. As discussed in Section 2.2, Bowyer and Kahne (2020) show that different digital civic learning experiences produce different civic outcomes, indicating that digital political learning depends on specific forms of engagement. Generative AI matters because, unlike earlier digital media, it does not simply distribute or filter information; it actively produces text, arguments, and simulated deliberation. It therefore becomes a potential intermediary in civic knowledge formation. Research on political learning in online settings further suggests that digital discussion environments shape political knowledge, meaning that once conversational AI enters this space, the mechanisms of political socialization also shift. At the same time, political socialization in digital environments is shaped not only by exposure but also by credibility judgments and interpretive authority. Generative AI is analytically distinctive because its fluent and authoritative-seeming responses may amplify risks such as misinformation, overconfidence, and reduced source evaluation while also enabling new forms of inquiry. In this study, therefore, generative AI is treated as a socio-technical mediator that influences how students acquire information, test positions, and form civic dispositions. This extension supports the article's focus on civic and political education rather than higher education more generally.

3.2. Analytical Framework: The "Technology-Pedagogy-Governance" Model

Building on these foundations, this study develops a Technology-Pedagogy-Governance (TPG) model to capture how generative AI becomes educationally meaningful in higher education. The model rests on a simple premise: AI cannot be understood through tools alone. Its effects emerge through the interaction of technological configuration, pedagogical design, and governance arrangements. The first dimension, technology, refers to the AI tools, systems, and modes of deployment used by institutions. It asks not only what systems are adopted, but also whether AI functions as an optional assistant, embedded learning resource, public platform, or institutional service. In line with digital governance

theory, tools are understood as operating within institutional architectures rather than acting independently. The second dimension, pedagogy, concerns how AI affects teaching design and learning outcomes. The central issue is not access to AI, but whether it enriches inquiry, reflection, and critical reasoning, or instead displaces them. Grounded in experiential learning theory, this dimension evaluates whether AI supports the purposes of civic and political education, including argument evaluation, source criticism, ethical reflection, and deliberation. The third dimension, governance, includes the policies, ethical safeguards, and oversight mechanisms regulating AI use. Drawing on digital governance theory and the ethics literature reviewed in Section 2.3, governance covers assessment rules, disclosure requirements, integrity policies, and risk management. This dimension is essential because, without clear frameworks, the same technology may produce inconsistent and educationally harmful outcomes. UCL's categorization of AI use in assessment is a clear example of governance aligning risk management with pedagogical intent. The TPG model is therefore relational rather than additive. Technology without pedagogy leads to instrumental adoption; pedagogy without governance produces inconsistency; and governance without the other two yields formal compliance without educational coherence. This structure also justifies the comparative case design, since the selected institutions differ not simply by context but by how they configure these three dimensions. The model guides the case analysis in Chapter 5 and underpins the transferable synthesis in Chapter 6.

Table 2. Analytical Dimensions for Case Study

Dimension	Core Focus	Analytical Questions
Technology	AI tools and systems	What AI is used and how?
Pedagogy	Teaching design and learning outcomes	How does AI shape learning?
Governance	Ethics, policy, oversight	How are risks managed?

Table 2 operationalizes the TPG model for the empirical analysis. In Chapter 4, it informs case selection and data collection; in Chapter 5, it provides a common structure for comparing Arizona State University, University College London, and the University of Helsinki.

3.3. Criteria for Case Selection

Case selection follows a logic of purposeful, theory-informed comparison rather than statistical representativeness. In line with Helström's account of multiple-case design, cases were chosen because each illuminates a distinct configuration of the TPG dimensions while remaining comparable within higher education. The aim was not to identify the "best" universities, but to select analytically strong and empirically traceable examples of institutional responses to generative AI. Three criteria guided selection. First, each case had to demonstrate a clearly documented institutional response to generative AI or AI literacy. This ensured evidentiary traceability and minimized speculative interpretation. Arizona State University met this criterion through its OpenAI collaboration; University College London through

its publicly available assessment framework; and the University of Helsinki through its long-running Elements of AI initiative. Second, the cases had to represent meaningful variation across technology, pedagogy, and governance. Arizona State University represents institutional deployment and experimentation, UCL represents governance and assessment alignment, and Helsinki represents scalable AI literacy. Together, these cases provide theoretically productive contrast across the TPG model. Third, the cases had to offer plausible transferability for emerging universities. This follows directly from the study's aim in Section 1.3. While Arizona State University's enterprise-level model may be difficult to replicate, it remains analytically useful as an example of organized institutional experimentation. UCL's governance model is more readily adaptable because assessment and integrity rules can be translated into less resource-intensive settings. Helsinki's model is especially transferable because its open-access format demonstrates how AI literacy can precede or accompany more advanced implementation. Taken together, these criteria ensure that case selection is neither arbitrary nor prestige-driven. Each case offers documented practice, meaningful variation, and lessons relevant to institutions seeking responsible and pedagogically defensible approaches to generative AI. These criteria also prepare the transition to Chapter 4, which details the methodological design and analytical procedures.

4. Methodology

4.1. Research Design: Multiple Case Study

This study adopts a qualitative multiple-case study design to examine how higher education institutions respond to generative artificial intelligence in domains related to civic learning, public reasoning, and responsible digital education. Following Helström (2026), case study research is appropriate when a contemporary phenomenon must be examined in its real-life context and when the boundary between phenomenon and context is not clearly defined. Generative AI in higher education meets these conditions: its significance cannot be separated from the institutional settings in which it is governed, taught, and normalized. The research therefore focuses on institutional practice rather than abstract technological capacity. A multiple-case design was selected for both analytical and substantive reasons. It enables structured comparison across institutions that differ in mission, context, and AI adoption, thereby strengthening explanatory power. It is also aligned with the comparative research questions in Section 1.4, which ask not only how generative AI is used, but also which pedagogical and governance arrangements appear effective and transferable. Consistent with the TPG model developed in Chapter 3, each case is treated as both a bounded setting and a unit of analytical contrast. The study is instrumental rather than descriptive. It does not aim to provide full institutional histories of Harvard University, University College London, or the University of Helsinki. Instead, each case is used to illuminate how

universities organize the relationship among AI tools, pedagogical design, and governance safeguards in educational contexts relevant to civic competence and responsible participation. This directly addresses the gap identified in Section 2.5 regarding the lack of practice-oriented, transferable institutional models. A further strength of the design is that it supports analytic rather than statistical generalization. The purpose is not to represent all universities, but to use selected cases to refine and test the TPG framework. The methodology is therefore theory-led and empirically grounded: the cases make visible different institutional pathways of AI adoption, while the framework provides a common basis for comparison.

4.2. Case Selection Strategy

The three cases were selected through purposeful, theory-informed sampling. The aim was not to choose the most prominent universities using AI, but to identify institutions that show meaningful variation across the TPG dimensions-technology, pedagogy, and governance-while remaining sufficiently documented for transparent qualitative analysis. As outlined in Section 3.3, selection was guided by three criteria: documented institutional response, meaningful variation in institutional focus, and partial transferability to less resourced settings.

4.2.1. Harvard University (United States): AI-assisted civic education

Harvard University was selected because it provides evidence of AI-assisted educational practice in settings linked to civic learning, public policy, and democratic inquiry, while also maintaining institution-wide guidance on responsible use. Harvard's Office of Undergraduate Education and related units have issued guidance for instructors and students covering syllabus policy, privacy, copyright, and academic integrity. Harvard Kennedy School strengthens the case by providing a concrete example of AI use in civic-oriented teaching. Liz McKenna described using Code Interpreter in a course on social movements to support data visualization, enabling students without programming backgrounds to engage in public-interest analysis (McKenna, 2025). Harvard is therefore relevant not because it has a single centralized "AI and civic education" program, but because it combines institutional guidance with traceable course-level experimentation.

4.2.2. University College London (United Kingdom): AI governance in teaching

University College London (UCL) was selected because it offers one of the clearest public examples of institutional governance for generative AI in teaching and assessment. UCL distinguishes among prohibited, assistive, and integral uses of AI in assessment, thereby operationalizing governance in a pedagogically differentiated manner. Rather than treating AI as wholly legitimate or illegitimate, UCL links acceptable use to assessment purpose, academic integrity, and educational design. UCL is also supported by both official guidance and emerging scholarly reflection. Public materials describe generative AI as something to be used "effectively and

ethically," while publications in UCL Discovery discuss strategies for assessment integration. This combination makes UCL the strongest case for examining governance within the TPG framework.

4.2.3. University of Helsinki (Finland): AI literacy and ethics education

The University of Helsinki was selected because it represents AI literacy and ethics-oriented public education at scale. Its Elements of AI initiative has become one of the most visible AI literacy programs globally. By May 2023, the course had reached more than one million learners in 170 countries and had been translated into 26 languages. This makes Helsinki important not only because of scale, but because it frames AI education as broad civic competence rather than specialist training. The case is further strengthened by Helsinki's wider open online offerings, including courses on AI and democracy, justice and security, discrimination, and ethics. In contrast to Harvard's course-linked experimentation and UCL's governance-centered model, Helsinki foregrounds a pedagogical strategy in which ethical and civic understanding is built through open, scalable literacy education. It therefore provides the clearest case for examining pedagogy and transferability within the TPG framework. Taken together, the three cases constitute a maximum-relevance comparative sample. Harvard highlights AI-assisted civic and public-policy learning, UCL provides a governance-centered model, and Helsinki demonstrates scalable AI literacy and ethics education. They were selected not to represent national systems, but to capture three analytically distinct institutional responses to generative AI in higher education.

4.3. Data Collection

Data collection was based entirely on publicly accessible documentary sources. Because the study investigates institutional models rather than individual experiences, documentary evidence provides the most appropriate primary source base. It allows direct analysis of how universities define, justify, and structure AI use. Four categories of material were collected for each case: institutional reports, policy documents, academic publications, and open-access course materials. Institutional reports included official announcements, strategic communications, and summaries of initiatives. These materials reveal institutional priorities, legitimation strategies, and intended uses. Policy documents were especially important for UCL and Harvard, as they show how institutions translate ethical concerns into practical rules on assessment, acceptable use, and integrity. Academic publications were used mainly to contextualize and interpret documentary evidence, particularly where official guidance was concise. Open-access course materials were especially valuable for Harvard and Helsinki because they offered direct insight into pedagogical design and AI integration in teaching. To strengthen credibility, the study used source triangulation within each case. Major interpretive claims were not based on a single source type when corroborating evidence was available. Harvard's inclusion, for

example, depended on both general guidance and course-level evidence; UCL's governance model was supported by policy documents and repository publications; Helsinki's AI literacy orientation was supported by both official reporting and course-platform materials. This reduced the risk of overinterpreting promotional or isolated documents.

4.4. Data Analysis Method: Thematic Analysis

The collected materials were analyzed through thematic analysis, following Braun and Clarke (2006). This method was selected because it is flexible and well suited to heterogeneous qualitative materials, including policy documents, institutional webpages, course descriptions, and scholarly texts. It enabled the identification of recurrent patterns across cases while retaining sensitivity to institutional differences. The analysis proceeded in four stages. First, the documentary corpus for each institution was read repeatedly to identify material relevant to the three TPG dimensions introduced in Section 3.2. Second, the materials were coded both deductively and inductively. Deductive coding used the TPG framework-technology, pedagogy, and governance-while inductive coding identified emergent subthemes such as disclosure, AI literacy, academic integrity, accessibility, public reasoning, and scalability. Third, codes were compared within and across cases to identify broader themes and case-specific configurations. Fourth, these themes were refined through iterative comparison with the theoretical framework in Chapter 3 and the literature reviewed in Chapter 2. This procedure enabled the study to move beyond descriptive cataloguing toward analysis of how institutions configure the relationship among AI tools, pedagogy, and governance. For instance, "explicit assessment categorization" was particularly salient in the UCL case, "open-access literacy at scale" in the Helsinki case, and "course-level experimentation in civic or public-policy teaching" in the Harvard case. To enhance rigor, three principles were followed. First, coding decisions were grounded in publicly available textual evidence rather than inferred intention. Second, themes were retained only when supported by repeated or especially clear documentation. Third, the analysis remained attentive to "negative space," including issues not prominently addressed in institutional materials, such as transferability limits or the absence of direct student outcome evidence. Official documents were therefore treated as evidence of institutional positioning, not proof of effectiveness.

Table 3. Overview of Selected Case Institutions

Institution	Country	Focus Area	Data Sources
Harvard University	USA	AI in civic engagement education	Course design materials, faculty guidance, institutional AI guidance pages
University College London	UK	AI governance in higher education	Policy frameworks, assessment guidance, university repository publications
University of Helsinki	Finland	AI literacy programs	Open online courses, institutional reporting, course platform materials

Table 3 summarizes the empirical basis of the study. The cases are comparable not

because they are identical, but because each provides a documented institutional configuration of the technology, pedagogy, and governance dimensions outlined in Table 2.

5. Case Findings and Cross-Case Analysis

This chapter presents findings from the three cases introduced in Chapter 4 and interprets them through the Technology-Pedagogy-Governance (TPG) model developed in Chapter 3. Sections 5.1-5.3 present within-case findings for Harvard University, University College London, and the University of Helsinki, while Sections 5.4-5.5 identify cross-case opportunities and challenges. In line with the research questions in Section 1.4, the focus is not on whether generative AI is present in higher education, but on how it is used, governed, and pedagogically framed in ways that may inform transferable institutional learning.

5.1. Case 1: Harvard University (USA)

Harvard University combines institutional guidance, course-level experimentation, and public-policy-oriented AI teaching. University guidance frames generative AI as pedagogically usable while emphasizing privacy, security, copyright, and course-specific rules. At Harvard Kennedy School (HKS), this is complemented by course offerings focused on generative AI and public leadership. The case is therefore best understood as a layered institutional setting in which AI is introduced under controlled academic discretion. A key finding is that Harvard's AI-related teaching emphasizes civic reasoning rather than technical proficiency. The HKS course The Science and Implications of Generative AI explores how AI works, how it can be used effectively, and how it may shape society, while explicitly addressing risks and responsible use. This orientation places AI within the domain of public judgment and institutional responsibility, directly linking it to the civic-learning concerns outlined in Sections 1.2 and 2.2. Harvard also provides evidence of AI-supported civic reasoning and debate-oriented learning. Liz McKenna's use of Code Interpreter in a course on social movements shows how AI can lower technical barriers and enable students without programming backgrounds to engage in public-interest analysis. In TPG terms, Harvard's technology dimension is characterized by AI as a teaching assistant and analytical scaffold, rather than a replacement for substantive reasoning. Pedagogically, Harvard links AI use with argumentation, reflection, and critical thinking. Teaching resources encourage using AI to support learning design, while Bok Center materials explicitly frame AI as both a learning partner and an object of critical inquiry. This aligns with the experiential learning perspective in Section 3.1.2: AI is valuable when it extends inquiry rather than bypassing it. Harvard's strength thus lies in structured augmentation-expanding students' analytical reach while keeping human judgment central. At the same time, Harvard explicitly recognizes the risk of over-reliance on AI-generated perspectives. HKS rules state that, unless permitted by instructors,

AI-generated content in coursework violates the academic code. Harvard's guidance similarly insists that acceptable use must be clearly defined. The case therefore illustrates both the promise and fragility of AI-supported civic education: generative AI can widen access to policy-relevant analysis, but it also risks substituting machine fluency for independent judgment.

5.2. Case 2: University College London (UK)

Where Harvard foregrounds pedagogical experimentation, University College London (UCL) is defined by institutional governance of AI in education. UCL has developed a public framework distinguishing among prohibited, assistive, and integral uses of generative AI in assessment. This is not a peripheral advisory tool but the core of UCL's response. In TPG terms, UCL is the strongest governance-centered case in the sample. The most significant finding is that governance at UCL is used to structure pedagogy rather than merely police misconduct. Staff are expected to clarify acceptable AI use in relation to specific assessments, meaning that governance is directly tied to learning outcomes and task design. This supports the argument in Section 3.2 that governance without pedagogy is insufficient: at UCL, policy helps make teaching design coherent in an AI-rich environment. A second finding concerns ethical framing. UCL's student-facing guidance presents AI use as something to be undertaken "effectively and ethically," not simply efficiently. The same logic appears in its academic-integrity materials, which stress that AI permissibility depends on the way it is used and the category of assessment involved. This avoids binary governance and instead treats AI regulation as contextual, rule-based, and instructional. UCL also frames academic integrity as an assessment design problem, not merely a student compliance issue. The article Integrating Generative AI in Academic Assessments identifies UCL's three-tiered framework as a mechanism for reconciling integrity with the reality of AI use. This makes UCL the clearest example of how universities can preserve standards without denying the presence of generative AI in academic life. Its main weakness, however, is implementation complexity. A nuanced governance framework requires staff to interpret categories consistently, communicate them clearly, and apply them across diverse assessment types. UCL therefore demonstrates a central tension in AI governance: the more sophisticated and pedagogically responsive the framework, the more demanding its implementation becomes.

5.3. Case 3: University of Helsinki (Finland)

The University of Helsinki represents a different institutional pathway: AI literacy as a foundation for civic competence. Its most visible initiative is Elements of AI, a free online course created to introduce non-specialists to the basics of artificial intelligence. By May 2023, the course had reached more than one million learners in 170 countries. This indicates a model built around broad access and public

educational reach rather than elite institutional infrastructure. Helsinki's relevance lies in the fact that its AI education model extends beyond technical literacy to include societal and ethical understanding. Its open online course catalogue includes not only Elements of AI and Building AI, but also courses on AI and democracy, justice and security, discrimination, and ethics. In TPG terms, Helsinki's pedagogical dimension is defined by critical awareness rather than passive use. A central finding is that Helsinki treats AI literacy as a precondition for meaningful participation in an AI-mediated society. Rather than beginning with efficiency or institutional deployment, its courses frame understanding AI's assumptions, limitations, and social consequences as an educational good in itself. This aligns closely with the theoretical discussion in Section 3.1.3, where AI was conceptualized as a mediator of public knowledge and civic judgment. The major strength of the Helsinki model is public accessibility and scalability. Because its courses are open and designed for broad participation, they offer a potentially transferable model for institutions that lack enterprise-level AI infrastructure. Its limitation is that, compared with Harvard or UCL, it provides less evidence of deep, discipline-specific integration into course-level assessment or debate-intensive teaching. Helsinki is therefore strongest as a model for critical preparation rather than advanced classroom deployment.

Table 4. Cross-Case Comparison of AI Applications

Dimension	Harvard	UCL	Helsinki
Core Focus	Civic engagement	Governance	AI literacy
AI Usage	Teaching assistant	Policy-regulated tool	Educational content
Strength	Critical thinking	Ethical control	Public accessibility
Weakness	Risk of dependency	Implementation complexity	Limited depth in application

Table 4 shows that the three universities represent distinct institutional emphases rather than competing versions of the same model: Harvard prioritizes pedagogically guided civic use, UCL governance and integrity, and Helsinki literacy, ethics, and accessibility.

5.4. Key Opportunities Identified

The comparison reveals three major opportunities.

First, enhanced engagement emerges most clearly from Harvard. HKS materials suggest that generative AI can lower technical barriers, support exploratory analysis, and enrich discussion in courses tied to public policy and social movements. When used as an assistant rather than an answer-generator, AI can strengthen student engagement with contested public issues. Second, ethical awareness through structured governance is most evident in the UCL case. Its three-category framework and integrity guidance show that governance can prevent uncritical use while also functioning as an educational signal: students learn that AI use is contextual, rule-bound, and ethically consequential. Third, scalable education

through open-access AI literacy is best illustrated by Helsinki. The reach of Elements of AI and the wider MOOC ecosystem demonstrates that universities can build educational capacity around AI without relying on proprietary infrastructure. For emerging universities, this model is valuable not only because of scale but because it suggests a feasible starting point.

5.5. Key Challenges Identified

The comparison also identifies three recurring challenges. The first is algorithmic bias and the epistemic limits of AI-generated content, which appears across all three cases. Harvard's emphasis on responsible use, UCL's category-specific guidance, and Helsinki's ethics-oriented courses all indicate that AI outputs cannot be treated as neutral or automatically reliable. This finding is consistent with the concerns raised in Sections 2.1 and 2.3 by Safdar et al. (2026) and Mochizuki et al. (2025). The second is over-reliance, most visible in the Harvard case. The greater the role of AI in reasoning, writing, or analysis, the greater the risk that students will rely on machine-generated fluency instead of developing independent arguments. Harvard's explicit limits on unauthorized AI use show institutional recognition of this problem. The third is governance complexity, most evident in the UCL case. Nuanced frameworks are educationally preferable to blanket prohibitions, but they require staff training, communication, monitoring, and interpretive consistency. For less resourced institutions, the challenge may therefore lie not in adopting governance in principle, but in implementing it effectively. Taken together, these findings answer the research questions posed in Section 1.4. Internationally, generative AI is being used in higher education in at least three ways relevant to civic and political learning: as a pedagogical assistant in public-facing analytical contexts, as a governed educational tool embedded in assessment regimes, and as a literacy object through which students and wider publics learn to understand AI's societal implications. The evidence also suggests that effective institutional responses depend not on technology alone, but on its pedagogical and governance framing. The most transferable lesson is that universities need not begin with advanced deployment; more feasible pathways may start with literacy, governance, and carefully bounded pedagogical experimentation. This conclusion provides the basis for the transferable model developed in Chapter 6.

6. Discussion and Transferable Model

6.1. Synthesis of Findings

The cross-case analysis shows that the educational significance of generative AI in higher education lies not in the technology itself, but in the institutional configuration through which it is introduced. Across the three cases, effective use depends on aligning tool design, pedagogical purpose, and governance architecture. Harvard shows that AI can support civic reasoning when embedded in policy-oriented tasks

and framed as a means of experimentation, analysis, and informed judgment rather than as a source of finished answers. Harvard Kennedy School's public materials emphasize responsible use for societal benefit and highlight tasks requiring critical thinking and human-based learning.

UCL, by contrast, demonstrates that governance is not secondary to adoption but constitutive of educational legitimacy. Its three-category framework-prohibited, assistive, and integral use-links AI regulation directly to assessment design and learning outcomes. UCL's student guidance further stresses that AI permissibility depends on context and on whether use creates unfair advantage, making governance more nuanced than blanket prohibition. The Helsinki case shows that AI literacy can function as civic preparation in its own right. Elements of AI has reached more than one million learners in 170 countries and has been translated into 26 languages, while Helsinki's wider course ecosystem addresses democracy, justice, discrimination, and ethics. Together, these materials suggest that Helsinki's model is less about classroom deployment than about building the capacity to interpret AI critically before or alongside application. Taken together, the three cases suggest that the most transferable lesson is sequencing. Universities need not begin with enterprise-scale deployment. A more defensible pathway begins with literacy, proceeds through bounded pedagogical experimentation, and is stabilized through explicit governance. This conclusion is consistent with the broader literature, which shows that the value of AI in higher education depends on design rather than novelty alone (Crompton & Burke, 2023; Safdar et al., 2026). Across cases, AI adds value only when it remains subordinate to educational purpose and human judgment.

6.2. Construction of a Transferable Model

On the basis of these findings, this study proposes a Transferable AI-Enhanced Civic Education Model organized around the three dimensions established in Chapter 3: technology, pedagogy, and governance. The model is transferable not because it can be copied directly across settings, but because it identifies components that can be adapted sequentially and scaled according to institutional capacity. Its core premise is that, in civic and politically relevant education, generative AI should extend inquiry, reflection, and analytical participation rather than replace them. The first component is controlled AI integration. Harvard shows that AI can support public-policy learning while preserving student responsibility for intellectual work. In practice, this means using AI for brainstorming, simulation, low-stakes feedback, structured data exploration, or scenario testing, but not as a substitute for independent argumentation or evidence evaluation. The second component is critical thinking-centered pedagogy. Across the cases, the strongest educational uses of AI were those that encouraged students to test prompts, compare outputs, detect limitations, and reflect on quality. Harvard's teaching guidance emphasizes creativity and practical application, while Helsinki's literacy-and-ethics ecosystem frames AI as something to be critically examined. In civic education, this implies assignments that

require students to compare positions, justify judgments, interrogate generated claims, and connect AI outputs to primary sources and human discussion. The third component is clear ethical governance. UCL provides the clearest operational model. Its three-category framework and associated integrity guidance show how institutions can distinguish among prohibited, assistive, and integral uses while preserving fairness and making expectations explicit. The transferable lesson is not that every university should replicate UCL's exact taxonomy, but that viable adoption requires formal guidance, disclosure norms, assessment alignment, and locally intelligible rules. This is especially important in civic education, where hallucination, bias, and algorithmic error can distort public reasoning as well as academic quality.

Table 5. Transferable "AI-Enhanced Civic Education Model"

Component	Strategy	Practical Implication
Technology	Controlled AI integration	Use AI as assistant, not replacement
Pedagogy	Critical thinking-centered design	Encourage reflection and debate
Governance	Clear ethical guidelines	Institutional regulation needed

Table 5 distills the cross-case findings into a model that is minimal in structure but robust in principle. It does not assume access to proprietary systems or advanced technical capacity. Instead, it identifies the minimum conditions under which AI-enhanced civic education is likely to remain educationally legitimate and practically workable.

6.3. Policy and Institutional Implications

The first implication is that universities should build AI-supported but human-led educational systems. None of the cases supports a fully automated vision of higher education. Even the most innovation-oriented example, Harvard, frames responsible use and informed judgment as central goals. Institutional strategy should therefore prioritize augmentation over substitution. The second implication is that universities need institutional AI governance frameworks, not merely ad hoc classroom decisions. UCL demonstrates that when expectations are made explicit and linked to assessment categories, governance becomes educative as well as regulatory. At minimum, policy should address acceptable use, disclosure, academic integrity, and staff responsibility for communicating assessment-specific rules. The third implication is that universities should promote AI literacy before large-scale AI application. Helsinki shows that foundational understanding of AI can be built through open and accessible formats before institutions invest in extensive deployment. For less resourced universities, this sequence is likely to be more sustainable and less risky than beginning with proprietary integration.

6.4. Practical Recommendations for Emerging Universities

For emerging universities, the first recommendation is to start with pilot courses rather than institution-wide transformation. Pilot courses allow institutions to test educational value, gather feedback, and refine guidance before scaling. The Harvard

case suggests that hands-on, policy-relevant, and discussion-oriented modules offer productive entry points. The second recommendation is to use open resources wherever possible, especially in early stages. The Helsinki model shows that AI literacy and ethics education can be developed through open online materials with broad reach. Adapting existing public resources lowers entry costs while establishing a common conceptual baseline among staff and students. The third recommendation is to develop internal AI guidelines early, even if initially simple. UCL's approach shows that clarity about acceptable use is a precondition for coherent adoption. Emerging universities may not need a fully elaborated three-category model immediately, but they do need clear rules on when AI is prohibited, assistive, or required, together with basic expectations for acknowledgment and integrity.

6.5. Limitations

This study has two main limitations. First, it relies on secondary and documentary data. The analysis reconstructs institutional positioning and practice models from reports, policy documents, and course materials, but it cannot directly verify how consistently these models operate in everyday teaching or how students experience them. Official documentation therefore provides evidence of institutional positioning rather than demonstrated effectiveness. Second, the study is based on a small number of cases. Harvard, UCL, and Helsinki were selected for analytical contrast and transferability, not statistical representativeness. They are useful for theory-informed comparison but cannot capture the full range of institutional responses to generative AI in higher education.

6.6. Future Research Directions

Future research should extend this analysis in two directions. The first is quantitative validation. The model proposed here could be tested through survey-based or mixed-method research to examine whether institutions combining bounded AI use, critical pedagogy, and explicit governance produce better outcomes in civic reasoning, ethical awareness, or academic integrity. The second is longitudinal impact research. Because generative AI is evolving rapidly, it remains unclear how stable current institutional arrangements will be over time. Longitudinal studies could examine whether literacy-first approaches produce more resilient adaptation, whether governance frameworks remain effective as tools improve, and whether AI-supported civic learning generates durable changes in students' reasoning habits.

Acknowledgements

This research was supported by the Student Affairs Research Project of Huizhou University (2025), entitled "Opportunities and Challenges of Generative Artificial Intelligence in Youth Political Guidance: Exploring and Practicing the Empowerment of the 'Grand Ideological and Political Education' Framework", for which the author served as the principal investigator. The project was approved as a Key Project

following institutional review and evaluation.

References

[1] Alam, A. (2026). Pedagogies in a Digitally Mediated World for the Future of Civic Education: Educating Youth for Hashtag Citizenship. In Youth Civic Engagement and Political Participation in Digital Spaces (pp. 55-92). IGI Global Scientific Publishing.

[2] Bezeriță, L. (2025). SOME PECULIARITIES AND PERSPECTIVES OF THE FINNISH EDUCATION SYSTEM. Internal Auditing & Risk Management, 71(1), 77-85.

[3] Bo, N. S. W. (2025). OECD digital education outlook 2023: Towards an effective education ecosystem. Hungarian Educational Research Journal, 15(2), 284-289.

[4] Bowyer, B., & Kahne, J. (2020). The digital dimensions of civic education: Assessing the effects of learning opportunities. Journal of Applied Developmental Psychology, 69, 101162.

[5] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative research in psychology, 3(2), 77-101.

[6] 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.

[7] Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: the state of the field. International journal of educational technology in higher education, 20(1), 1-22.

[8] Gupta, V., & Nyamapfene, A. (2025). Generative AI in universities: Practices at UCL and other institutions, and the path forward. Internet Reference Services Quarterly, 29(1), 131-151.

[9] Helström, K. (2026). Case Study Methodology in Qualitative Research. In Approaches to Qualitative Research (pp. 141-147). Routledge.

[10] Kusumawardani, V., Agustin, D. A. C., & Permatasari, S. S. (2025). Social Media and Political Preferences Among Generation Z in the 2024 Indonesian General Election. Mediakita, 9(2), 246-269.

[11] Lee, T. S. C. (2025). Magic lines: Artist-teachers' preschool art studio curriculum design and implementation. In Arts and creativity in East and Southeast Asian preschools: Towards 'glocal'pedagogies (pp. 139-154). Singapore: Springer Nature Singapore.

[12] Liu, R., & Glave, C. Z. (2026). Justice, Equity, Diversity, and Inclusion in Artificial Intelligence and Childhood Education: A Critical Systematic Review. Journal of Research in Childhood Education, 40(1), 5-23.

[13] McKenna, T. J. (2025). Making Sense of Sensemaking: Designing Authentic K-12 STEM Learning Experiences. Teachers College Press.

[14] Mochizuki, Y., Bruillard, E., & Bryan, A. (2025). The ethics of AI or techno-solutionism? UNESCO's policy guidance on AI in education. British Journal of Sociology of Education, 1-22.

[15] Mott, R., & Haddad, B. (2025). A call to better Qual: A philosophical and methodological examination to advance case study research. Journal of Agricultural Education, 66(1), 41-41.

[16] Nguyen, H. (2025). Bringing AI into teaching: Understanding Vietnamese teachers' perspectives and pedagogical challenges. International Journal of Educational Methodology, 11(3), 335-347.

[17] Safdar, H., Hashmi, K., & Shah, M. A. (2026). Systematic Literature Review of the Governance of Integrity, Equity, Privacy, and Accountability in the Use of Artificial Intelligence Models in University Classrooms: Ethical Principles. Research Consortium Archive, 4(1), 311-333.

[18] Shaikh, S., Maroof, M. A., Kaskar, P., Girhe, P. B., Magar, A. V., & Sudhahar, W. (2026). Generative AI and Ethical Integration in K-12 and Higher Education: Policies, Pedagogy,

and Practice. In Teacher Perspectives and Responsible Practice for Integrating AI in the Classroom (pp. 89-118). IGI Global Scientific Publishing.

[19] Tlili, A. (2025). Between Fear and Desire, the Monster Artificial Intelligence (AI): Analysis through the Lenses of Monster Theory. arXiv preprint arXiv:2508.08318.