Journal of Systemics, Cybernetics and Informatics (2025) 23(4), 51-56
https://doi.org/10.54808/JSCI.23.04.51

Reflexivity as a Compass: The European AI Act and Its Implications for U.S. Higher Education Institutions

Jasmin COWIN
Richard P. Nathan Public Policy Fellow
Rockefeller Institute of Government
Albany, NY, 12203

ABSTRACT

This narrative analysis explores how the European Union's Artificial Intelligence Act (EU AI Act) holds the potential to shape institutional discourse of U.S. higher education. As artificial intelligence becomes deeply embedded in university operations, from admissions and instruction to monitoring and assessment, it raises urgent questions about institutional purpose, power, and accountability. Drawing on Kantian ethics, the analysis highlights the tension between external regulatory structures and internal moral reasoning. The EU AI Act (Regulation 2024/1689), with its risk-based classification of AI systems and its extraterritorial provisions, introduces binding obligations for transparency, oversight, and ethical alignment in educational applications. These obligations challenge existing norms of voluntary governance in U.S. academia and signal a shift toward anticipatory and structured forms of technological oversight. Within this landscape, reflexivity is positioned not as a rhetorical gesture but as a necessary institutional capacity. It refers to the ongoing process of self-examination that engages with embedded assumptions, power dynamics, and the normative dimensions of algorithmic systems. This analysis argues that reflexivity must guide institutional responses to AI governance if universities are to align technological adoption with their academic values and global responsibilities.

Keywords: reflexivity, artificial intelligence, higher education, EU AI Act, ethical governance, socio-technical systems, transdisciplinarity.

1. INTRODUCTION

"Two things fill the mind with ever new and increasing admiration and awe the more often and more enduringly reflection is occupied with them: the starry heavens above me and the moral law within me." Kant (1998, p. 5)

As Immanuel Kant in his *Critique of Pure Reason* recognized, awe arises not only from the external cosmos but also from the ethical dimensions of our internal reasoning. This tension between external systems and internal moral agency is particularly significant today, as artificial intelligence (AI) becomes embedded in the institutional fabric of higher education. The rapid adoption of AI across universities has reshaped core functions from admissions decisions and student engagement to research methodologies and administrative operations. These changes extend far beyond questions of technical implementation. They raise foundational concerns about power, epistemology, and institutional purpose. As AI systems become more integral to educational decision-making, they bring with them both transformative opportunities and structural risks. For this reason, critical engagement with how such systems are designed, deployed, and governed is now becoming a central focus of higher education institutions.

2. THE EUROPEAN AI ACT: A PARADIGM SHIFT IN TECHNOLOGICAL GOVERNANCE

The EU AI Act, published in the Official Journal of the European Union on July 12, 2024, provides a comprehensive and binding legal framework to regulate the development and deployment of artificial intelligence across the EU. The Act recognizes that "AI is a fast evolving family of technologies that contributes to a wide array of economic, environmental and societal benefits across the entire spectrum of industries and social activities" (Regulation (EU) 2024/1689, 2024, Recital 4). It establishes a graduated model based on risk that categorizes AI systems by their potential threat to fundamental rights and safety, assigning legal obligations accordingly. Of particular relevance to higher education institutions, Annex III of the Act specifically designates several educational applications as AI systems with high risk, including:

- AI systems intended to determine access or admission to educational institutions at all levels
- AI systems intended to evaluate learning outcomes, including when those outcomes are used to steer the learning process
- AI systems intended to assess the appropriate level of education that an individual will receive or be able to access
- AI systems intended for monitoring and detecting prohibited behavior of students during tests (Regulation (EU) 2024/1689, Annex III, Section 3)

While the primary jurisdiction of the Act remains the EU, its extraterritorial impact is both real and imminent. U.S. higher education institutions that engage in research across borders, host EU students, or utilize AI systems developed within the EU are now entangled in an emerging compliance landscape. These institutions, long accustomed to voluntary governance and academic freedom, must now respond to externally imposed accountability structures. Within this shifting terrain, reflexivity is not simply a theoretical position but a necessary compass for policy interpretation, institutional response, and epistemic recalibration.

Although the EU AI Act does not explicitly refer to reflexivity, its requirements underscore the importance of institutional self-awareness and ethical accountability. The Act introduces a tiered classification of AI systems based on risk, with high risk as those used in education, subject to strict obligations for transparency, documentation, human oversight, and impact assessment (European Commission, 2024). These requirements signal a

broader shift from permissive innovation to anticipatory governance, requiring universities to interrogate the ethical and social implications of their AI infrastructures.

This regulatory turn is not limited to Europe. As U.S.-based edtech companies and higher education institutions increasingly engage with international students or operate within the European digital space, they are also subject to the Act's extraterritorial provisions (GovTech, 2024; KPMG Law, 2025). Moreover, the European AI Act is likely to influence U.S. policy indirectly by shaping expectations around risk management, human-centered design, and algorithmic accountability (Brookings Institution, 2023). In this shifting policy landscape, reflexivity becomes essential, not merely as a rhetorical or theoretical construct but as a practical and organizational capacity.

Reflexivity, in this context, refers to the ability of institutions to critically examine their own practices, assumptions, and frameworks for decision making. It goes beyond retrospective reflection by requiring an ongoing, embedded process of self-questioning. This process includes attention to the power relations that shape institutional action, the assumptions embedded in algorithmic systems, and the normative commitments underlying educational goals (Mäntymäki et al., 2022). Reflexivity is therefore not an ancillary concern; it is a compass that can help universities navigate tensions between technological possibility and academic purpose, between institutional efficiency and human values.

3. LITERATURE REVIEW

"For pure reason is never related directly to objects, but instead to concepts of them given by the understanding." Kant (1998, p. 406)

Kant's insight that "pure reason is never related directly to objects, but instead to concepts of them given by the understanding" identifies a structural constraint that remains relevant in present-day AI governance. The EU AI Act reflects this same epistemic limitation. Like pure reason, the Act does not engage AI systems directly. Instead, it operates through preconstructed categories that define what counts as risk, trust, or transparency. These categories are necessary to regulate complex systems, but they also rely on abstractions that are removed from the concrete realities they aim to manage.

Laux, Wachter, and Mittelstadt (2023) argue that the Act's treatment of trust exemplifies this problem. They critique the assumption that trustworthiness can be established through pre-market conformity assessments. In their view, this collapses the difference between regulatory risk acceptability and the ongoing process of building and sustaining trust. Trust, they argue, is not a fixed property of a system but a condition that requires continuous oversight, communicative transparency, and responsiveness to context. Their critique draws attention to the Act's reliance on procedural assurance rather than substantive institutional accountability. In educational settings, where AI systems interact with vulnerable groups and influence long-term life paths, this substitution is not merely technical. It has ethical and political consequences.

Veale and Zuiderveen Borgesius (2021) examine the structure of the AI Act in relation to the European Union's history of product safety regulation. The Act applies a four-part classification of risk: unacceptable, high, limited, and minimal. This framework is designed to support both innovation and fundamental rights. However, the authors question the applicability of product regulation models to adaptive systems like machine learning, which develop behavior through data and use, rather than fixed design. Regulatory categories, although necessary for administrative action, remain abstract. They simplify complexity rather than engage with it. These categories function as what Kant would call "concepts given by the understanding": cognitive structures that organize experience, but do not access the object itself. In this case, the object is a learning system whose behavior changes over time. The classification scheme brings the system into a governance framework, but only by setting aside many of its defining properties.

The case of educational AI systems makes this abstraction explicit. As outlined in Annex III of the AI Act, systems used in university admissions, assessment, and student monitoring are designated as high risk. This classification carries obligations for human oversight, transparency, documentation, and impact assessment (European Commission, 2024). Hauer et al. (2023) analyzed the application of these rules to AI tools in German education. They found that only about thirty systems would currently qualify as high risk, but they also stressed that the process of classification is itself unstable. Determining whether a system falls within a particular category requires interpretation. It depends not only on technical features, but also on institutional priorities, legal judgments, and the framing of use cases. Hauer and colleagues argue for adaptive compliance models that can adjust as both the regulation and the technologies evolve. Their findings show that even empirical assessments of the Act's scope rely on conceptual mediation. This supports Kant's broader claim. Regulatory reason, like pure reason, cannot bypass the role of conceptual structures. It does not interact with systems directly but through frameworks that organize them in advance. This becomes particularly visible in education, where AI systems influence what Hauer et al. (2023) call "long-term life trajectories." In such contexts, the stakes exceed operational accuracy. They extend into domains of practical reason, where questions of fairness, autonomy, and human development become central. When regulatory concepts fail to account for these dimensions, they do not simply miss technical details. They risk reinforcing patterns of exclusion or misrecognition under the appearance of procedural neutrality.

Further concerns arise from the Act's provisions on explainability. Chung et al. (2024) state that "There's a general consensus that increasing explainability will lead to transparent, reliable, and accountable AI, which contributes to greater safety." And further argue that these provisions risk generating what they describe as a "false sense of security." While the Act requires that high-risk systems be understandable to users, it does not define what interpretability means in technical or operational terms. This lack of specification leaves institutions and developers with wide discretion in determining compliance. As a result, systems may be presented as transparent even when their logic remains opaque. In education, where AI tools are used to allocate opportunity and assess performance, this misalignment can result in serious harm. Systems that appear understandable may in fact obscure how decisions are made, reducing the possibility of critique or contestation.

Taken together, these perspectives point to a recurring structural issue. The AI Act, like the forms of reason Kant described,

organizes its object of concern through abstract conceptual schemes. These schemes are necessary for action, but they cannot eliminate ambiguity, contestation, or interpretive instability. Effective governance of AI for educational institutions, therefore, requires more than technical compliance. It demands reflexivity. Institutions must examine how their own assumptions shape the ways they interpret, implement, and justify AI systems. Without this reflexive capacity, the abstractions that enable governance may also undermine it.

4. FROM KANTIAN EPISTEMOLOGY TO REFLEXIVE AI GOVERNANCE

The epistemic structure of the EU AI Act necessitates what we might call reflexive governance, a systematic approach that recognizes the mediated relationship between regulatory knowledge and the technological systems it seeks to regulate. Like Kant's notion of pure reason, which cannot directly access objects but must engage them through concepts constructed by the understanding, the AI Act operates through abstract categories rather than direct interaction with AI systems themselves. These categories help structure legal accountability, but they cannot fully grasp the changing nature of the systems they aim to oversee. The result is a gap between legal classification and technological behavior that cannot be closed through rule application alone. Bridging this gap requires reflective judgment, the ability to respond to particular cases when no general rule can be presupposed.

The EU AI Act exemplifies this challenge. With more than 50,000 words, 180 recitals, 113 articles, and 13 annexes, the Act sets out what it describes as "a uniform set of legal obligations for all actors involved in the development and use of AI." However, this comprehensiveness reveals an underlying assumption: that a sufficiently detailed categorical framework can provide epistemic and normative access to AI systems. The Act adopts a use-based rather than a technology-based approach. It applies rules to specific AI applications rather than to their underlying technical architecture. This allows for flexibility, but it also reflects a deeper limitation. No fixed regulatory category can fully describe systems whose properties emerge through training data, user interaction, and ongoing adaptation. The belief that exhaustive rule systems can fully determine future cases risks what might be described, borrowing from Kantian terminology, as a form of regulatory paralogism, a failure of judgment rooted in conceptual overreach.

5. REFLEXIVITY: THEORETICAL FOUNDATIONS AND PRACTICAL APPLICATIONS

Reflexivity is central to understanding how institutions can respond to the uncertainties inherent in AI governance. Its philosophical foundation appears in early modern thought, particularly in Descartes' turn toward self-reflection. It is developed systematically in Kant, whose critical project investigates not the objects of knowledge themselves but the conditions under which knowledge is possible. In the *Critique of Pure Reason*, Kant asserts that "The I think must be able to accompany all my representations; for otherwise something would be represented in me that could not be thought at all" (Kant, p. 246). This notion of transcendental apperception is reflexive in structure. It requires the subject to be aware of its own cognitive activity as the ground for unified experience.

Kant further distinguishes between determinative judgment and reflective judgment. Determinative judgment applies known rules to particular cases. Reflective judgment, by contrast, arises when a case is given but no rule is available in advance. In such situations, the task of judgment is to find or create the appropriate principle. As Kant writes, "If the universal (the rule, the principle, the law) is given, then the power of judgment, which subsumes the particular under it... is determining. If, however, only the particular is given and the universal has to be found for it, then the power of judgment is merely reflecting" (Kant, p. 67). Reflective judgment offers a model for reflexive practice under conditions of uncertainty. It is an enabler to assess and revise the categories they use to interpret novel situations.

This model is directly relevant to the AI Act's treatment of educational technologies. The Act classifies systems used for admissions, student evaluation, and monitoring as high risk. It imposes requirements for transparency, documentation, and human oversight. These requirements acknowledge that such systems cannot be governed by predetermined technical rules alone. Instead, their regulation requires ongoing assessment of how they affect fundamental rights and institutional responsibilities. In these cases, reflective judgment must operate not only at the level of system design but also at the level of institutional interpretation. Regulatory reason must be supplemented by institutional reflexivity, structured processes through which universities evaluate the effects of AI systems and adjust their governance practices accordingly.

6. THE CONSEQUENCES OF MODERNITY AND REFLEXIVITY

In sociology, Anthony Giddens (1990) conceptualized reflexivity as a defining characteristic of late modernity, where social practices are constantly examined and reformed in light of incoming information about those practices. In *The Consequences of Modernity*, Giddens argues that modernity is marked by a heightened level of reflexivity, meaning that both individuals and institutions are engaged in a continuous process of self-examination and adaptation. This reflexivity distinguishes modern societies from traditional ones, where social practices tend to be taken for granted and remain relatively unchanged (Giddens, 1990).

Giddens defines reflexivity as the ongoing monitoring and revision of social practices based on new information and knowledge. This process is not limited to individuals but is embedded within the very institutions of modernity, affecting everything from personal identity to global economic systems. Reflexivity is facilitated by the proliferation of abstract systems and expert knowledge, which provide the resources for individuals and organizations to navigate the complexities and uncertainties of modern life (Giddens, 1990).

The EU AI Act exemplifies this late modern reflexivity through what Bradford (2020) terms the "Brussels Effect," whereby EU regulations shape global standards through market mechanisms and normative influence. This extraterritorial reach creates what we might call "regulatory reflexivity," as institutions worldwide must examine and adapt their practices in response to EU standards, regardless of their jurisdictional location. U.S. higher education institutions with partnerships in the European Union, or those using AI-enhanced platforms developed in Europe, find themselves compelled into reflexive practices as they assess

whether they are subject to provisions in the Act. This recursive relationship between knowledge and action becomes particularly salient in AI contexts, where systems both shape and are shaped by the data they process and the decisions they inform.

Beck, Giddens, and Lash (1994) extend this analysis through their theory of reflexive modernization, arguing that contemporary societies must develop capacities for self-confrontation with the unintended consequences of technological progress. They distinguish between simple modernization, which applies rationalist principles to traditional society, and reflexive modernization, which turns critical examination back upon the modernization process itself. As Beck et al. note, "The reflexivity of modernity actually subverts modernity" (1994, p. 2), suggesting that genuine progress requires continuous questioning of foundational assumptions. The EU AI Act embodies this reflexive modernization by challenging traditional approaches to technology regulation. Rather than applying existing product safety frameworks directly to AI systems, the Act forces a fundamental reconsideration of regulatory purposes and methods. As the analysis reveals, "translating product safety principles to AI systems raises significant implementation questions" because "conformity assessment procedures, borrowed from traditional product regulation, may prove inadequate for systems that evolve through machine learning." This framework proves particularly relevant for understanding how universities must approach AI integration: not as a simple application of new tools to existing processes, but as a fundamental reconsideration of educational purposes and methods.

Bourdieu and Wacquant (1992) contribute a critical dimension to reflexivity through their emphasis on the social construction of knowledge and the need for researchers to examine their own position within fields of power. Their concept of "epistemic reflexivity" demands that scholars interrogate the conditions of possibility for their own knowledge production, including institutional pressures, disciplinary boundaries, and implicit assumptions. For universities implementing AI systems under the EU AI Act, this Bourdieusian reflexivity requires examining how institutional habitus shapes technology adoption, whose interests are served by particular implementations, and what forms of knowledge are privileged or marginalized through algorithmic mediation. The Act's requirements for "mandatory impact assessments, human oversight provisions, and transparency obligations" force institutions into precisely this kind of reflexive self-examination, questioning not just technical compliance but the educational and social implications of AI integration.

7. REFLEXIVITY INTEGRATION WITH AI GOVERNANCE

The intersection of reflexivity theory with practical AI governance reveals both synergies and tensions. Schiff et al. (2021) analyze how reflexive practices can be operationalized in AI development through participatory design methods, stakeholder engagement, and iterative evaluation processes. Their empirical studies of AI labs demonstrate that reflexivity often conflicts with commercial pressures for rapid deployment and competitive advantage, suggesting that regulatory frameworks like the EU AI Act may serve as necessary external prompts for reflexive practices that might otherwise be marginalized.

Educational contexts present unique challenges and opportunities for reflexive AI governance. Selwyn (2019) examines how educational technologies embody particular pedagogical assumptions and power relations, often reinforcing rather than challenging existing inequalities. His analysis suggests that reflexivity in educational AI must extend beyond technical considerations to examine fundamental questions about the purposes of education, the nature of learning, and the distribution of educational opportunities. This aligns with Porayska-Pomsta's (2024) critique of technocentric approaches in educational AI, which “prioritize algorithmic efficiency over pedagogical effectiveness and student agency.”

The transformative potential of reflexivity in educational settings is empirically demonstrated by Heijmans and Eweg (2023), who studied Living Labs as sites for developing what they term “transformative capacity.” They define this as “the ability to deal with diversity in norms, values, beliefs, encourage reflexivity when facilitating multi-stakeholder processes and collectively create knowledge that is actionable and contributes to sustainable change” (Heijmans and Eweg, p. 347). Their findings suggest that reflexivity cannot be mandated through regulation alone but must be cultivated through structured practices that bring together diverse perspectives and create space for fundamental questioning of assumptions.

For U.S. higher education institutions navigating the EU AI Act's requirements, the literature points toward reflexivity as both a conceptual framework and practical necessity. The Act's extraterritorial reach means that institutions cannot simply comply with its technical requirements but must develop capacities for ongoing self-examination and adaptation. This reflexive orientation offers a path between minimal compliance and paralytic over-caution, enabling institutions to harness AI's educational potential while maintaining critical awareness of its limitations and risks. As universities develop new governance frameworks spanning multiple jurisdictions and value systems, reflexivity emerges not as an abstract ideal but as an essential compass for navigating complexity.

Understanding reflexivity in this context requires disentangling it from mere reflection. Reflection often involves retrospective analysis of prior actions. Reflexivity, by contrast, involves situated awareness of the assumptions, power relations, and semiotic systems through which knowledge is constructed and decisions are made. As Esposito, Freda, and De Luca Picione (2021) write, reflexivity entails “a critical understanding of the symbolic and relational configurations that define the possibilities of knowledge” (p. 5). It becomes central when institutions face opaque regulatory demands, especially those articulated through a different legal and ethical culture. For U.S. academic institutions to engage meaningfully, they must recognize that legal compliance alone is insufficient. Instead, institutional actors need to interrogate their own epistemological orientation toward AI: how AI is conceptualized, implemented, and justified.

8. THE PARADOX OF ALGORITHMIC REFLEXIVITY

A Kantian lens exposes a deep paradox at the heart of reflexive AI governance in higher education. As Dreyfus (2001) argues in his critique of artificial intelligence, “What distinguishes persons from machines, no matter how cleverly constructed, is that a human being’s actions... are based on situational understanding”

(p. 47). Machine learning systems, however optimized, lack embodied judgment and cannot evaluate the assumptions that structure their operations. As a result, university deployments of AI for admissions, grading, or research evaluation under the EU AI Act may generate what Introna (2016) terms "reflexivity theaters," where institutional routines mimic critical reflection while concealing underlying algorithmic rigidity.

Still, this critique must be situated within a broader shift in transnational higher education that the Act is already shaping. Its regulatory logic extends beyond compliance, catalyzing changes in the institutional conditions of cross-border academic engagement. A clear example of this is the evolution of Erasmus+, the EU's flagship education program. While Erasmus+ remains committed to international cooperation and maintains partnerships beyond the EU, recent developments have shifted focus inward. The program increasingly emphasizes its European dimension and alignment with overarching EU policy goals, such as the European Education Area, the Digital Education Action Plan, and the European Skills Agenda. The program demonstrates how policy mechanisms can extend influence beyond their immediate regulatory targets, reshaping institutional priorities and international norms.

9. REGULATORY CONVERGENCE AND GRAVITATIONAL DRIFT

"There are, therefore, only three species of these dialectical syllogisms, as many as there are ideas in which their conclusions result. In the first class of syllogisms, from the transcendental concept of a subject that contains nothing manifold I infer the absolute unity of this subject itself, even though in this way I have no concept at all of it. This dialectical inference I will call a transcendental paralogism." Kant, I. (1998, p 416)

Kant's account of transcendental paralogisms offers a lens through which to understand the structural assumptions underlying regulatory models like the EU AI Act. The Act risks replicating a form of dialectical error: by relying on formal compliance structures to infer substantive trustworthiness, it treats categorical unity as epistemological and ethical certainty. The resulting framework imposes procedural clarity that may, in practice, obscure the very complexity it intends to govern.

In contrast, recent U.S. initiatives such as the Presidential AI Challenge present a bottom-up model that cultivates participation and practical problem-solving. Launched by the White House in 2025, it is structured to build a national AI talent pipeline while anchoring technological development in democratic education. It creates distinct tracks that validate critical thinking (Track I), empower builders and coders (Track II), and elevate educators as co-equal innovators (Track III). Its philosophical commitments are not regulatory but pedagogical. It is not simply a competition but a national effort to seed an AI-literate public capable of addressing real-world challenges. By challenging educators to design pedagogy itself, the initiative makes a profound statement: the most critical infrastructure for future technological leadership is not regulatory compliance, but the classroom that embodies innovation.

As Cowin (2025) demonstrates in her analysis of this initiative, there exists a fundamental philosophical divergence between American and European approaches to AI governance:

> "This distinctly American, bottom-up approach of fostering innovation through competition and inspiration stands in stark contrast to the path taken by the European Union. The EU's landmark AI Act operates from a top-down, regulatory philosophy. Where the Presidential Challenge acts as a catalyst, the AI Act functions as a set of guardrails. It is a comprehensive legal framework that categorizes AI systems by risk and imposes strict rules and transparency obligations, particularly on high-risk applications, with the primary goal of protecting fundamental rights and user safety. The two approaches reveal a fundamental divergence in strategy: the U.S. is prioritizing the releasing of potential through a national challenge, while Europe is prioritizing the containment of risk through preemptive legislation." (Cowin, 2025)

This contrast does not merely reflect different administrative tools. It signals divergent views on how societies cultivate technological futures. Cowin continues:

> "The initiative suggests that the path to a national 'Golden Age' is paved not just with raw code, but with the cultivation of virtuous citizens who know how to wield AI tools in service of supporting and elevating their communities and their country." (Cowin, 2025)

10. CONCLUSION: THE ARC OF TRANSDISCIPLINARITY

When American educational institutions and educators participate in innovation-driven programs and European institutions operate under the constraints of the EU AI Act, they engage with fundamentally different normative systems. These engagements can form transdisciplinary assemblages. These are constellations of inquiry, reflection, and institutional learning that do not require harmonization. They generate knowledge through structured tension rather than alignment. This transdisciplinary arc allows collaboration while preserving normative plurality. American institutions can continue to pursue innovation and the cultivation of virtue. European institutions can learn from these experiments without compromising their legal commitments to human dignity and proportionality.

As Nicolescu (2002) argues, transdisciplinarity occurs between disciplines, across disciplines, and beyond disciplinary boundaries. It is not a toolkit or policy instrument. It is a condition of practice under complexity. It requires institutions to remain within unresolved tensions, to sustain dialogue across incompatible norms, and to develop accountability that is not reducible to either innovation logics or regulatory procedure. Reflexivity, then, is reframed not as a predefined institutional posture but as a shared epistemic orientation that emerges from working within divergent systems.

ACKNOWLEDGEMENTS

I wish to acknowledge with deep appreciation the contribution and feedback of Prof. Dr. Birgit Oberer, Founder and Director at ETCOP and Secretary General at IGIP, who served as a non-anonymous reviewer for this article.

REFERENCES

[1] Beck, U., Giddens, A., & Lash, S. (1994). Reflexive modernization: Politics, tradition and aesthetics in the modern social order. Stanford University Press.

[2] Bourdieu, P., & Wacquant, L. (1992). An invitation to reflexive sociology. University of Chicago Press.

[3] Bradford, A. (2020). The Brussels effect: How the European Union rules the world. Oxford University Press.

[4] Brookings Institution. (2023, April 25). The EU and U.S. diverge on AI regulation: A transatlantic comparison and steps to alignment. https://www.brookings.edu/articles/the-eu-and-us-diverge-on-ai-regulation-a-transatlantic-comparison-and-steps-to-alignment/

[5] Chung, C., Nguyen, T., & Breuer, J. (2024). Explainability and the false sense of security in AI governance. Journal of AI Policy and Society, 12(1), 45–62. https://arxiv.org/html/2405.03820v1

[6] Cowin, J. (2025, July 31). To win AI race, White House enlists students and educators for Presidential AI Challenge. Stankevicius. Retrieved from https://stankevicius.co/artificial-intelligence/to-win-ai-race-white-house-enlists-students-and-educators-for-presidential-ai-challenge/

[7] Dreyfus, H. L. (2009). On the internet (2nd ed.). Routledge.

[8] Esposito, G., Freda, M. F., & De Luca Picione, R. (2021). Reflexivity: Applying a transdisciplinary approach to the study of human experience. Theory & Psychology, 31(5), 655–677.

[9] European Commission. (2024). Artificial Intelligence Act (Regulation EU 2024/1689). https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence

[10] European Commission. (2024). Artificial Intelligence Act (Regulation EU 2024/1689). https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence

[11] Giddens, A. (1990). The consequences of modernity. Cambridge: Polity Press.

[12] GovTech. (2024, March 26). What will the EU AI Act mean for edtech in the U.S.? https://www.govtech.com/education/higher-ed/what-will-the-eu-ai-act-mean-for-ed-tech-in-the-u-s

[13] Hauer, P., Krafft, T. D., Sesing-Wagenpfeil, J., & Zweig, K. A. (2023). Quantitative study about the estimated impact of the AI Act. European Journal of Technology Regulation, 8(2), 145–167.

[14] Heijmans, A., Eweg, R. (2023). Transformative research and education in Living Labs using the SDGs as a compass. International Journal of Sustainability in Higher Education, 24(2), 335–352.

[15] Introna, L. D. (2016). Algorithms, governance, and governmentality: On governing academic writing. Science, Technology, & Human Values, 41(1), 17–49.

[16] Kant, I. (1998). Critique of pure reason (P. Guyer & A. W. Wood, Eds. & Trans.). Cambridge University Press.

[17] Kant, I. (1998). Critique of pure reason (P. Guyer & A. W. Wood, Eds. & Trans.). Cambridge University Press.

[18] KPMG Law. (2025, March 20). AI Act: This applies to AI in universities and research. https://kpmg-law.de/en/ai-act-this-applies-to-ai-in-universities-and-research/

[19] Laux, J., Wachter, S., & Mittelstadt, B. (2022). Trustworthy artificial intelligence and the European Union AI Act: On the conflation of trustworthiness and the acceptability of risk. SSRN. https://doi.org/10.2139/ssrn.4230294

[20] Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Putting AI ethics into practice: The hourglass model of organizational AI governance. arXiv. https://arxiv.org/abs/2206.00335

[21] Nicolescu, B. (2002). Manifesto of transdisciplinarity. State University of New York Press.

[22] Porayska-Pomsta, K. (2024). From algorithm worship to the art of human learning: Insights from 50-year journey of AI in education. International Journal of Artificial Intelligence in Education, 34(1), 45–72. https://arxiv.org/abs/2403.05544

[23] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on laying down harmonised rules on artificial intelligence and amending certain Union legislative acts. Official Journal of the European Union. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

[24] Schiff, D., Rakova, B., Ayesh, A., Fanti, A., & Lennon, M. (2021). Principles to practices for responsible AI: Closing the gap. arXiv preprint arXiv:2006.04707.

[25] Selwyn, N. (2019). Should robots replace teachers? AI and the future of education. Polity Press.

[26] Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU artificial intelligence act. Computer Law Review International, 22(4), 97–112.