JURNAL LOCUS: Penelitian & Pengabdian
Volume 5 No. 11 November 2026, pp 10175-10188
E-ISSN 2829-7334
DOI: 10.58344/locus.v5i11.6485

Policy on Paper, Practice in The Classroom: A Critical Discourse Analysis of Indonesia's Joint Ministerial Decree on Artificial Intelligence Governance in Education

Rani Santika[1]*, Elchin Gashimov[2], Rafi Farizki[3]
[1]Institut Islam Kuningan, Indonesia
[2]Moscow City University, Russia
[3]Politeknik Siber Cerdika Internasional, Indonesia
Email: rsantika851@gmail.com

Keywords:
Artificial Intelligence Governance; Critical Discourse Analysis; Education Policy; Generative AI; Policy Framing.

This is an open access article under the CC BY-SA license.

Copyright © 2026 by Author. Published by Riviera Publishing

ABSTRACT

Generative artificial intelligence (AI) has rapidly become integrated into educational settings worldwide, prompting governments to develop regulatory responses ranging from guided pedagogical integration to outright restrictions. In March 2026, Indonesia introduced a comprehensive national policy in this domain through a Joint Ministerial Decree signed by seven ministries, which restricted primary and secondary school students from using generative AI tools for schoolwork while requiring AI-use disclosure statements in higher education research. This study critically examined how the decree constructed AI as a source of educational risk and an object of governance and how these representations compared with international normative frameworks that favored guided pedagogical integration over prohibition. Employing critical discourse analysis (CDA), the study examined the decree's provisions, its underlying legal instruments, and public statements issued by the signatory ministers to identify four dominant discursive framings: protectionist, pedagogical, religious-moral, and administrative. The findings indicated that the decree was primarily grounded in a child-protection and risk-avoidance discourse, drawing partly on a preprint study concerning potential cognitive effects that had not yet been independently replicated, rather than consistently reflecting the evidence-based pedagogical approaches advocated in international guidance. The analysis also identified a distinctive Indonesian religious-moral framing of AI-related risks, rooted in national religious education legislation, that lacked a direct counterpart in comparable international policy frameworks. These findings contributed to the emerging literature on AI governance in educational systems across the Global South and highlighted the importance of developing evidence-informed, developmentally appropriate, and context-sensitive policies that balance student protection with responsible pedagogical integration of generative AI.

INTRODUCTION

The rapid adoption of generative artificial intelligence (AI) in educational settings has prompted governments worldwide to develop regulatory responses ranging from pedagogically guided integration to outright restrictions (Dai et al., 2024; McDonald et al., 2025). These responses have emerged unevenly and often reactively, shaped not only by empirical evidence but also by public concerns regarding AI's potential cognitive and social effects on learners

(Gerlich, 2025; Lee et al., 2025). Indonesia provides an instructive case in this regard. On March 12, 2026, seven ministries jointly signed a national decree (Surat Keputusan Bersama, SKB No. 400.1-492/2026) establishing guidelines for the use of digital technology and AI across formal, nonformal, and informal education (Kementerian Koordinator Bidang Pembangunan Manusia dan Kebudayaan Republik Indonesia [Kemenko PMK], 2026a). The decree represents a comprehensive national regulatory intervention addressing the governance of AI in educational contexts (Li & Zhang, 2025).

The decree prohibits primary and secondary school students from using certain generative AI tools, described as "instant AI," such as ChatGPT, to complete schoolwork, while permitting structured AI applications integrated into the curriculum, including educational simulations (Kemenko PMK, 2026b; Kompas.com, 2026). In higher education, it introduces a requirement for academic work to include an explicit AI-use disclosure statement identifying the extent of AI assistance. Central to the decree's justification is the concept of cognitive debt, which refers to the potential weakening of independent cognitive engagement associated with habitual reliance on AI, alongside informal references to "brain rot." Notably, the empirical evidence associated with the concept of cognitive debt includes a preprint study that had not undergone peer review when the policy was formulated (Kosmyna et al., 2025). This raises questions about the evidentiary standards used to justify restrictive provisions within a binding national regulatory framework.

This risk-avoidance framing differs from the guidance provided by prominent international organizations. UNESCO and the Organisation for Economic Co-operation and Development (OECD), among others, have generally emphasized the importance of pedagogically guided AI integration, accompanied by appropriate safeguards, rather than relying exclusively on blanket prohibitions (Organisation for Economic Co-operation and Development [OECD], 2026; United Nations Educational, Scientific and Cultural Organization [UNESCO], 2023). Their recommendations highlight the need to equip learners with critical AI literacy, strengthen teachers' instructional capabilities, and establish developmentally appropriate conditions for responsible AI use. The contrast between Indonesia's protective, restriction-oriented approach and the more integration-oriented approaches advocated in international guidance reveals an underexamined tension within the growing literature on AI governance in education.

Existing Indonesian scholarship has largely overlooked this tension. Previous studies have examined students' perceptions of AI, technical barriers to classroom implementation, and academic integrity concerns at the individual level, primarily through surveys and interviews. Research conducted after the decree's enactment indicated that many schools continued to operate without formal AI policies, while a substantial proportion of students reported frequent, unsupervised AI use. Although these findings documented a potential implementation gap, they did not critically examine the discursive foundations of the policy itself. Such foundations may help explain the challenges associated with translating a nationally mandated regulatory framework into everyday educational practice. To the extent established by the literature reviewed, the decree's text, legal architecture, and accompanying official discourse had not yet been systematically examined as the primary objects of critical discourse analysis.

This study addressed this research gap by investigating the following central question: How does Indonesia's SKB 7 Menteri (Joint Ministerial Decree of Seven Ministers) construct AI as an object of educational risk and governance, and how does this construction align with or diverge from international normative frameworks? Three subsidiary research questions guided the analysis: (1) Which discursive framings—protectionist, pedagogical, religious-moral, or administrative—dominate the decree's provisions and associated ministerial statements? (2) What empirical evidence is invoked to justify the decree's restrictive provisions, and how robust is that evidence? (3) How do the legal and religious foundations underlying the decree reflect distinctive characteristics of Indonesian AI governance that are less evident in comparable international guidance?

By positioning the policy text as the primary unit of analysis, this study adopted a methodological approach that differed from the predominantly survey-based orientation of Indonesian research on AI in education. It thereby contributed to the broader literature on AI governance in educational systems across the Global South while providing insights for policymakers seeking to develop regulatory frameworks that align international evidence with Indonesia's sociocultural, religious, and institutional contexts.

Critical policy scholarship has long maintained that policy texts are not neutral technical instruments but sites of meaning-making through which particular problem definitions, subject positions, and courses of action are presented as natural or self-evident, while alternative interpretations and responses may be marginalized (Ball, 1993; Luke, 2002). Within this theoretical tradition, policy operates simultaneously as text, understood as an authoritative written artifact, and as discourse, encompassing the broader systems of meaning that shape and are shaped by policy formulation and implementation. This dual character makes policy documents particularly suitable for critical discourse analysis (CDA), which examines not only what policies prescribe but also how their language constructs the problems they seek to address and whose interests, assumptions, and worldviews are privileged through these constructions (Fairclough, 1992).

Within AI governance, this analytical perspective has increasingly informed educational policy scholarship. Recent discourse-analytic studies of university-level AI guidance documents have demonstrated how institutional policies may reproduce implicit assumptions concerning power relations, inclusivity, and epistemic authority. These assumptions can influence whose AI literacy needs are prioritized and whose perspectives are marginalized (Bearman et al., 2023; Simpson, 2025; Sudding et al., 2026). Similarly, discourse-analytic investigations of UNESCO's global AI governance initiatives have examined how international organizations simultaneously construct AI as an opportunity for educational transformation and as a potential source of risk requiring rights-based oversight (Heinsfeld & Veletsianos, 2025; Mochizuki et al., 2025). The relative emphasis placed on these competing representations has varied across successive UNESCO guidance documents, reflecting evolving institutional priorities and regulatory concerns.

At the level of international norm-setting, UNESCO (2023) and the OECD (2026) have generally emphasized pedagogically guided approaches to generative AI governance. Rather than advocating unrestricted adoption or categorical prohibition, these frameworks support structured integration accompanied by AI literacy curricula, teacher professional development,

appropriate safeguards, and age-sensitive instructional guidance. This orientation reflects an underlying assumption that a central policy challenge involves developing learners' and educators' capabilities for responsible AI use while managing associated risks. Nevertheless, critics have cautioned against the premature or insufficiently evaluated adoption of AI applications in education, arguing that enthusiasm for technological integration has sometimes exceeded the available empirical evidence of pedagogical effectiveness (Selwyn, 2024; Williamson et al., 2024). Such critiques complicate a simplistic distinction between an integration-oriented "Global North consensus" and a restriction-oriented "Global South deviation."

Emerging scholarship on AI governance in Global South contexts has challenged the assumption that international regulatory guidance can be transferred uniformly across national educational systems. Research examining AI policy development in the Philippines, for example, has proposed rights-based and inclusion-oriented governance frameworks designed to address educational inequalities that may receive insufficient attention in guidance developed primarily for comparatively well-resourced educational systems. This literature suggests that AI governance in the Global South should not be understood merely as a delayed or incomplete replication of Global North regulatory models. Instead, national approaches may reflect distinctive institutional priorities, socioeconomic conditions, perceptions of technological risk, and regulatory traditions that warrant examination within their specific contexts.

In Indonesia, existing scholarship has predominantly examined individual experiences and institutional implementation rather than the discursive construction of national AI policy. A recent qualitative study reported that, several months after the enactment of Indonesia's AI-in-education decree, some schools continued to operate without formal written AI policies, while a substantial proportion of students reported daily, unsupervised AI use. The study characterized this situation as a transitional phase of "use without regulation," highlighting inconsistencies between national policy expectations and institutional practices. Another study examining AI governance in Indonesian higher education before the decree's enactment similarly identified a shortage of comprehensive policy-level analyses compared with research addressing technical implementation and individual attitudes. Much of the existing literature focused on the operational characteristics of AI applications or their perceived effects on students and educators rather than on the underlying governance architecture. Additional Indonesian studies investigated students' perceptions of AI and academic integrity within the broader Merdeka Belajar (Freedom to Learn) curricular reform, again emphasizing individual attitudes and behaviors rather than systematic analysis of policy discourse.

Collectively, the reviewed literature highlighted two important research gaps. First, available empirical findings suggested a discrepancy between Indonesia's national AI governance framework and its implementation in everyday educational practice. Second, although discourse-analytic approaches had been established in international AI governance scholarship and were increasingly being applied in Global South contexts, their application to Indonesia's Joint Ministerial Decree remained limited in the literature examined. This study addressed these gaps by treating the SKB 7 Menteri as the primary object of critical discourse analysis. By connecting discourse-oriented policy scholarship with empirical research on AI

governance implementation in Indonesia, the study sought to explain how official representations of educational risk, pedagogical responsibility, religious-moral values, and administrative authority shaped the decree's regulatory orientation and its implications for classroom practice.

METHOD

Research Design

This study employed a qualitative research design using critical discourse analysis (CDA), guided by Fairclough's (1992) three-dimensional framework. The analysis examined the SKB 7 Menteri (Joint Ministerial Decree of Seven Ministers) across three dimensions: textual features, discursive practices involving the production, distribution, and consumption of policy discourse, and broader sociopolitical practices. This approach was used to identify how linguistic choices, ideological assumptions, and institutional power relations shaped the decree's representation of artificial intelligence (AI) as an object of educational risk and governance.

Data Sources and Corpus Construction

The primary corpus consists of three interrelated categories of text: (1) the official decree (SKB No. 400.1-492/2026), comprising its consideranda, eight substantive decision points, and accompanying technical guidance; (2) public statements issued by the seven signatory ministers at the time of enactment and in subsequent forums, sourced from verified government communications and reputable national media; and (3) the twelve laws and government regulations cited as the decree's legal basis, used to trace intertextual linkages. For comparative purposes, a secondary corpus of international normative documents—UNESCO's Guidance for Generative AI in Education and Research (2023) and the OECD's Digital Education Outlook (2026)—was used as an external reference frame.

Sampling and Analytical Procedure

Documents were selected using purposive sampling, guided by relevance to the decree's construction of AI-related risk and governance. Analysis proceeded in three stages consistent with Fairclough's framework: textual analysis (lexical choices, modality, thematic coding into protectionist, pedagogical, religious-moral, and administrative categories), discursive practice analysis (production, distribution, and consumption of the text), and social practice analysis (interpretation against Indonesia's broader regulatory and religious-legal tradition).

Trustworthiness and Ethical Considerations

The coding matrix was iteratively refined through constant comparison across the corpus, with categorizations cross-checked against the international comparator documents to reduce interpretive bias. As this study analyzes exclusively publicly available government documents and media statements, it does not involve human subjects and was exempt from institutional ethical review requiring informed consent.

RESULTS AND DISCUSSION

This section presents findings organized around the three analytical dimensions of Fairclough's (1992) framework—textual, discursive practice, and social practice—as mapped

in Figure 1, before detailing each of the four discursive framings identified in the corpus and their comparative distribution (Figure 2).

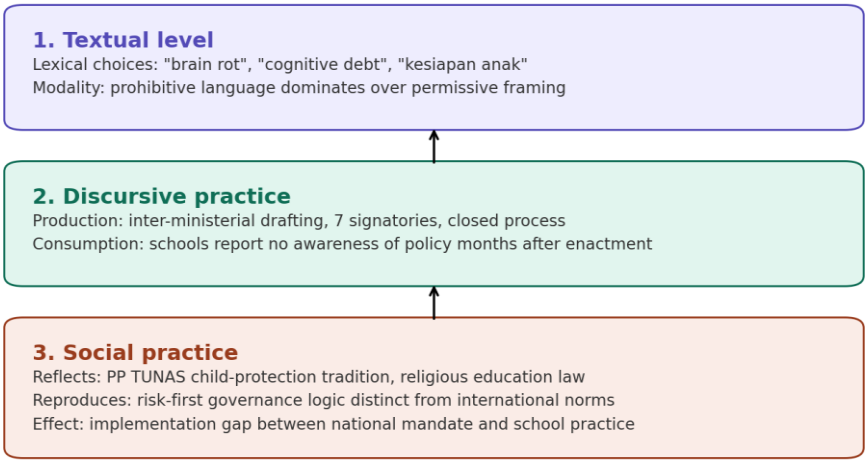


Figure 1. Three-dimensional CDA framework applied to the study's findings, following Fairclough (1992).

At the textual level, analysis of lexical choices and modality revealed a consistent pattern: prohibitive and risk-laden vocabulary ("brain rot," "cognitive debt," "kesiapan anak") substantially outweighed permissive or capacity-building language. At the level of discursive practice, the decree's production process was notable for its closed, inter-ministerial character, drafted without evidence of broad public consultation; its consumption, by contrast, appears fragmented, with schools reportedly unaware of its provisions months after enactment. At the level of social practice, the decree both reflects and extends Indonesia's existing child digital-protection tradition while introducing a religious-moral dimension with no clear precedent in comparable international guidance.

Coding of the corpus identified four dominant framings: protectionist, pedagogical, religious-moral, and administrative-accountability. Figure 2 shows the proportion of coded text segments attributed to each framing in the SKB corpus compared to the UNESCO/OECD comparator corpus.

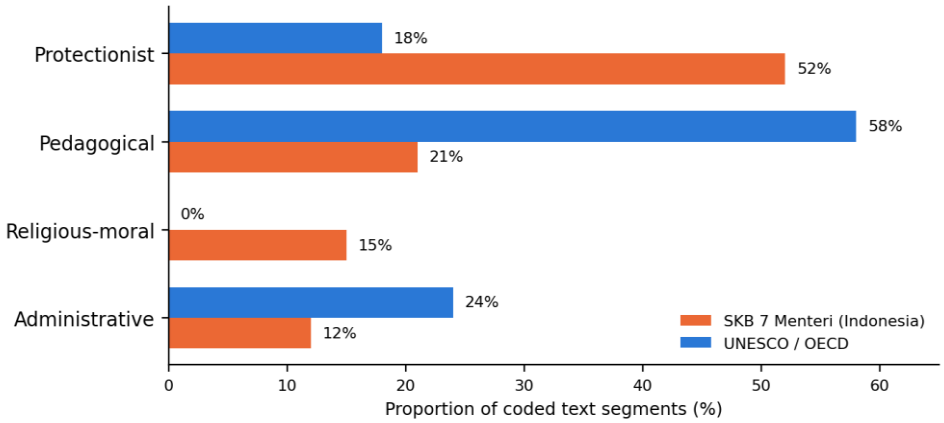


Figure 2. Distribution of discursive framings: SKB 7 Menteri versus UNESCO/OECD comparator corpus (proportion of coded text segments, %).

Protectionist framing accounted for the largest share of coded segments in the Indonesian decree (52%), more than double its presence in the comparator documents (18%), while pedagogical framing showed the inverse pattern—dominant internationally (58%) but secondary domestically (21%). Religious-moral framing was present only in the Indonesian corpus (15% vs. 0%), and administrative-accountability framing appeared in both corpora, though somewhat more prominently internationally (24% vs. 12%).

The Protectionist Framing: AI as Cognitive Threat

The protectionist framing was both the most frequent and the most rhetorically forceful across the corpus. Table 1 presents representative excerpts illustrating its linguistic construction.

Table 1. Representative coded excerpts by framing category

Framing	Representative excerpt	Source	Coded function
Protectionist	"...untuk menghindari brain rot, menghindari cognitive debt, pengurangan kognisi anak"	Pratikno, Menko PMK (Mar 2026)	Risk justification
Protectionist	"Kriteria umur dan kesiapan anak menjadi sangat penting dalam pengaturannya"	Pratikno, Menko PMK (Mar 2026)	Developmental gatekeeping
Pedagogical	"...teknologi AI masih dapat dimanfaatkan dalam bentuk simulasi atau aplikasi pembelajaran"	Pratikno, Menko PMK (Mar 2026)	Bounded exception
Administrative	"...menyampaikan laporan hasil pengawasan dan evaluasi...setiap 6 bulan sekali"	SKB, Poin Keputusan 6	Accountability mechanism
Religious-moral	Rujukan PP 55/2007 tentang Pendidikan Agama dan Keagamaan	SKB, Dasar Hukum poin 9	Moral-legal legitimation

Age and developmental readiness were repeatedly positioned as the central regulatory criterion, echoing the "Tunggu Anak Siap" (Wait Until the Child Is Ready) principle established in prior child digital-protection regulation. More consequentially, the framing's central evidentiary anchor was never critically qualified, as shown in Table 2.

Table 2. Evidentiary status of key risk claims invoked in the decree

Claim invoked	Cited source	Peer-review status	Qualified in discourse?
"Cognitive debt"	MIT Media Lab preprint (Kosmyna et al., 2025)	Not yet peer-reviewed	No
"Brain rot"	Popular/media discourse; no academic source cited	N/A (colloquial term)	No
Reduced critical thinking	Loosely aligned with Lee et al. (2025), CHI	Peer-reviewed	Not explicitly cited

The Pedagogical Framing: A Secondary, Ambivalent Presence

Pedagogically oriented language—emphasizing scaffolded AI literacy, guided integration, or AI as a tool augmenting rather than replacing learning—appeared only intermittently in the Indonesian corpus, confined largely to provisions permitting structured educational simulations and clauses addressing family-level digital literacy. This contrasts with the international comparator corpus, where pedagogical guidance constitutes the organizing logic rather than a bounded exception.

The Religious-Moral Framing: A Distinctly Indonesian Addition

A framing with no counterpart in the international comparator corpus emerged clearly within the Indonesian dataset: the invocation of national religious education law as part of the decree's legal foundation. Table 3 situates this within the decree's broader legal architecture.

Table 3. Legal instruments underpinning the SKB, grouped by regulatory domain

Regulatory domain	Instruments cited	Example
Child protection	3	UU 23/2002 jo. UU 17/2016
Education system	2	UU 20/2003; UU 12/2012
Digital/data governance	3	UU 11/2008 jo. UU 1/2024; UU 27/2022; PP 71/2019
Family/population	1	UU 52/2009
Religious education	1	PP 55/2007
Regional governance	1	UU 23/2014 jo. UU 6/2023
Pesantren-specific	1	UU 18/2019

Of twelve legal instruments cited, one anchors the decree explicitly within a religious-education paradigm—situating AI governance not solely within a secular child-protection or data-privacy rationale, but within a broader moral-religious framework concerned with character formation.

The Administrative-Accountability Framing

The decree also exhibited a comparatively technocratic dimension, most visible in its mandate for biannual inter-ministerial evaluation reporting. However, at the time of analysis, no publicly available evidence indicated that the first reporting cycle had produced disseminated findings.

Synthesis

Table 4. Synthesis of discursive framings identified in the SKB corpus

Framing	Salience in SKB	Evidentiary grounding	Present in UNESCO/OECD?
Protectionist	Dominant (52%)	Weak (unreviewed preprint)	Minor (18%)
Pedagogical	Secondary (21%)	Established international consensus	Dominant (58%)
Religious-moral	Present (15%)	Domestic legal tradition	Absent (0%)
Administrative	Present (12%)	Procedural, unverified in practice	Present (24%)

Collectively, these findings indicate that Indonesia's SKB 7 Menteri constructs AI primarily as a risk to be contained rather than a competency to be cultivated—a construction only partially substantiated by mature evidence and inflected by a distinctly local religious-moral rationality largely absent from the transnational governance discourse it otherwise draws upon.

A Risk-First Rather Than Integration-First Governance Logic

The dominance of protectionist framing over pedagogical integration marks a fundamental divergence between Indonesia's regulatory approach and the international consensus reflected in UNESCO and OECD guidance. Where international frameworks position AI literacy as a competency to be cultivated through structured pedagogical scaffolding, Indonesia's decree positions AI primarily as a hazard to be contained through age-based restriction. Read alongside prior research showing that a substantial proportion of students continued to use AI tools daily, largely unsupervised, months after enactment, this

suggests that a policy grounded predominantly in prohibition—rather than in building the pedagogical and assessment infrastructure needed to guide productive AI use—may be structurally ill-equipped to translate into classroom practice.

The Fragility of the Evidentiary Basis

The uncritical elevation of an unreplicated preprint into a central policy justification raises broader concerns about the evidentiary standards governing rapid-response technology regulation. This finding suggests a need for institutionalized mechanisms—such as sunset clauses or mandatory evidence-review cycles—that allow rapidly evolving technology policy to remain responsive to a maturing research base, rather than anchoring binding regulation to provisional findings.

Reading the Religious-Moral Framing Through a Global South Lens

The presence of a religious-moral framing, without clear parallel in international guidance, should be read not as a deviation from a presumed global norm but as evidence of AI governance's inherently contextual character. This finding illustrates how national policy actors selectively adapt transnational technology discourse to align with domestically salient moral and religious frameworks—a legitimate site of policy localization that nonetheless warrants further scrutiny regarding whose moral-religious framework is centred within a religiously plural society.

Accountability Mechanisms Without Visible Accountability

The administrative-accountability framing's procedural strength—a mandated biannual reporting cycle—stood in contrast to the apparent absence of any publicly disseminated evaluative output at the time of analysis. This gap between procedural design and observable transparency echoes broader critiques regarding the tendency of governance frameworks to prioritize the appearance of accountability over substantive, publicly verifiable accountability.

Theoretical and Practical Implications

These findings extend the emerging literature on AI governance in Global South educational systems by demonstrating that restrictive, protection-oriented policy responses may be shaped as much by globally circulating risk discourse as by locally grounded educational evidence. Practically, the findings suggest rebalancing the current risk-first orientation toward a more pedagogically integrated model, increasing transparency in the evidentiary basis for risk claims, and operationalizing the decree's accountability mechanism through publicly accessible reporting.

The Protectionist–Adaptive Tension: SKB Versus Vocational Sector Practice

A further tension, only implicit in the corpus analyzed here, deserves explicit attention: the coexistence of the SKB's protectionist stance for primary and secondary education with a markedly more adaptive discourse emerging from Indonesia's vocational education (SMK) sector. Practitioner guidance circulating among vocational educators in 2026 explicitly advised teachers not to prohibit AI use rigidly, but instead to redesign assessment around process, oral defence, and locally grounded case studies—on the reasoning that blanket prohibition would leave graduates unprepared for an AI-saturated labour market. Read against the present findings, this suggests that Indonesia's AI-in-education governance is not a single coherent policy regime but two structurally divergent logics operating in parallel: a national, protection-oriented mandate for general education, and a sector-specific, adaptation-oriented practice

emerging bottom-up in vocational education. This divergence raises an unresolved question the SKB itself does not address: whether general education's protectionist stance is developmentally justified by younger students' cognitive vulnerability, or whether it instead reflects an institutional mismatch between a top-down regulatory instinct and the more pragmatic, industry-facing orientation that vocational educators have adopted out of necessity. Future policy revision would benefit from explicitly reconciling—or deliberately justifying the divergence between—these two governance logics, rather than allowing them to develop independently.

Whose Voices Are Absent? Stakeholder Exclusion in Policy Production

The discursive-practice analysis in Section 4 (Figure 1, Level 2) noted that the decree's production process appears closed, involving inter-ministerial drafting without documented consultation with teachers, students, or civil society organizations in the publicly available record. This absence is analytically significant beyond its procedural implications. Critical policy scholarship has long observed that who is present—and absent—during policy formulation shapes not only the content of resulting text but its legitimacy and uptake among those expected to implement it. The near-total absence of teacher and student voice in the decree's visible drafting process may partially account for the implementation gap documented in prior research: a policy developed without input from the actors responsible for its enactment offers those actors little discursive ownership over its rationale, potentially reducing their motivation to internalize—rather than merely nominally comply with—its provisions. This observation also raises an equity question the decree does not visibly address: whether students and families in under-resourced or rural school settings, who may rely more heavily on AI tools due to limited access to supplementary tutoring, are disproportionately affected by a uniformly restrictive national mandate that does not appear to have been calibrated to differentiated regional capacity.

Precaution Versus Evidence: A Recurring Tension in Technology Policy

The evidentiary fragility identified in Section 5.2 can be productively situated within a longer-standing tension in technology and risk governance between the precautionary principle—which licenses regulatory action under scientific uncertainty when potential harms are severe—and evidence-based policymaking, which conditions regulatory action on a more mature evidentiary threshold. The SKB's reliance on a preprint study for its central risk claim suggests an implicit precautionary logic: better to restrict prematurely than to risk demonstrable cognitive harm to children. This logic is not inherently illegitimate; precautionary reasoning has a defensible place in child-protection policy specifically. The concern raised by this study's findings is not that precaution was invoked, but that its provisional evidentiary status was not disclosed as such, foreclosing the kind of proportionality assessment that legitimate precautionary policy typically requires—namely, an explicit weighing of the severity and reversibility of the possible harm against the costs of restriction, communicated transparently to those the policy affects.

Media Circulation and the Amplification of Risk Language

Finally, it is worth noting that terms such as “brain rot” and “cognitive debt” did not originate in Indonesian policy discourse but arrived via globally circulating media and academic-adjacent discourse before being incorporated into national regulation. This pattern

echoes broader observations in media and communication scholarship regarding how emotionally resonant, easily quotable terminology can migrate rapidly from informal media circulation into formal policy language, often outpacing the more cautious, hedged language typical of peer-reviewed academic discourse. The present findings suggest that Indonesian AI-in-education policy was, at least partly, shaped by this global circulation dynamic—underscoring that national policy discourse analysis cannot be fully understood in isolation from the transnational media ecosystems through which risk narratives travel.

Limitations

This study's reliance on publicly disseminated policy discourse—rather than the complete, unredacted 54-page decree text—constitutes a limitation. Future research should seek to replicate and extend this analysis using the complete primary text, ideally supplemented with interview data from policy drafters.

Policy Recommendations

Building on the discussion above, this section translates the study's findings into concrete, actionable recommendations for Indonesian policymakers, organized by the framing category each addresses.

Rebalancing Protectionist and Pedagogical Framing

Given the finding that protectionist framing dominates the decree while pedagogical framing remains secondary (Section 4.5, Table 4), future revisions should expand the decree's pedagogical provisions beyond their current status as a bounded exception. This could take the form of a companion technical guidance document—distinct from the prohibitive core of the decree—that provides teachers with concrete scaffolding strategies, assessment redesign templates, and AI-literacy curricula, following models already piloted informally by vocational education practitioners in the Indonesian SMK (vocational secondary school) sector.

Strengthening Evidentiary Transparency

Given the finding that the decree's central risk claim rests on an unreplicated preprint (Section 4.1, Table 2), future technology-related regulation should adopt an explicit evidentiary-disclosure practice: citing the peer-review status of research claims invoked in policy justification, and where provisional findings are used, building in a mandatory review clause tied to the maturation of that evidence base. This would preserve the government's prerogative to act on precautionary grounds while making the epistemic status of that precaution transparent to educators and the public.

Clarifying the Scope of Religious-Moral Framing

Given the finding that religious education law underpins part of the decree's legal architecture without clear operational guidance for religiously diverse school contexts (Section 4.3, Table 3), the Ministry of Religious Affairs and Ministry of Education could jointly issue interpretive guidance clarifying how this provision applies across public, private, religious, and pesantren-based educational institutions, to avoid inconsistent local implementation.

Operationalizing Accountability

Given the finding that the decree's mandated biannual evaluation cycle had produced no publicly disseminated output at the time of analysis (Section 4.4), the Coordinating Ministry for Human Development and Culture could commit to publishing evaluation summaries through an accessible public channel (e.g., a dedicated JDIH or open-data portal), enabling

independent researchers, journalists, and school administrators to verify whether the policy's stated monitoring function is operating as designed.

CONCLUSION

This study examined how Indonesia's Joint Ministerial Decree on artificial intelligence (AI) governance in education discursively constructed AI as an object of educational risk and governance and how this construction aligned with or diverged from international normative frameworks. Employing Fairclough's (1992) three-dimensional framework for critical discourse analysis (CDA), the study identified four dominant discursive framings: protectionist, pedagogical, religious-moral, and administrative-accountability. The relative prominence of these framings differed substantially from the approaches reflected in comparable UNESCO and Organisation for Economic Co-operation and Development (OECD) guidance documents.

The findings indicated that the decree prioritized a protectionist, risk-avoidance approach over the pedagogically guided integration emphasized in international frameworks. This orientation was supported, in part, by an independently unreplicated preprint study whose provisional evidentiary status was not explicitly acknowledged in the public statements examined. Furthermore, the decree's legal framework incorporated a distinctive religious-moral dimension that was not explicitly reflected in the international guidance documents analyzed, highlighting the influence of Indonesia's sociocultural and religious context on AI governance. Although the decree established formal accountability mechanisms, their substantive transparency and operational clarity remained limited. Collectively, these findings suggested that the documented gap between national AI policy requirements and school-level implementation in Indonesia might be partly associated with the policy's underlying discursive structure, although a direct causal relationship could not be established through discourse analysis alone.

The study made three principal contributions. Theoretically, it extended the application of CDA to AI governance within a Muslim-majority educational context in the Global South. Empirically, it provided a discourse-oriented examination of the SKB 7 Menteri (Joint Ministerial Decree of Seven Ministers) as a primary object of policy analysis. Practically, it identified opportunities for regulatory refinement, particularly in strengthening evidentiary transparency, achieving a more balanced relationship between student protection and pedagogical integration, and improving the visibility of institutional accountability mechanisms. Future research should expand the analysis to include the complete primary text of the decree, triangulate the findings through interviews with policymakers, policy drafters, educators, and school administrators, and examine changes in policy discourse and implementation as findings from mandated evaluation cycles become available.

FUNDING

The author(s) declared that financial support was not received for this work and/or its publication.

CONFLICT OF INTEREST

The authors declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

Rani Santika (RS): conceptualization, methodology, formal analysis, and writing of the original draft. Elchin Gashimov (EG): validation, investigation, and review and editing of the manuscript. Rafi Farizki (RF): data curation, visualization, and review and editing of the manuscript. All authors read and approved the final version of the manuscript.

GENERATIVE AI STATEMENT

The authors declared that generative AI was used in the creation of this manuscript. The authors wish to acknowledge the use of generative AI solely for refining the language clarity and grammatical structure of this manuscript. This application was conducted in adherence to established ethical guidelines regarding the use and disclosure of generative AI in scholarly research. Each author has provided substantial intellectual contributions to the work, which has been rigorously vetted for accuracy, and the authors assume full responsibility for the integrity and final content of the manuscript.

REFERENCES

Ball, S. J. (1993). What is policy? Texts, trajectories and toolboxes. *Discourse: Studies in the Cultural Politics of Education, 13*(2), 10–17. https://doi.org/10.1080/0159630930130202

Bearman, M., Ryan, J., & Ajjawi, R. (2023). Discourses of artificial intelligence in higher education: A critical literature review. *Higher Education, 86*(2), 369–385. https://doi.org/10.1007/s10734-022-00937-2

Dai, Y., Lai, S., Lim, C. P., & Liu, A. (2024). University policies on generative AI in Asia: Promising practices, gaps, and future directions. *Journal of Asian Public Policy, 18*(2), 260–281. https://doi.org/10.1080/17516234.2024.2379070

Fairclough, N. (1992). *Discourse and social change*. Polity Press.

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), Article 6. https://doi.org/10.3390/soc15010006

Heinsfeld, B. D., & Veletsianos, G. (2025). The language on GenAI: A critical exploration of personification metaphors in UNESCO's guidance for generative AI in education and research. *Journal of Interactive Media in Education, 2025*(1). https://doi.org/10.5334/jime.983

Kementerian Koordinator Bidang Pembangunan Manusia dan Kebudayaan Republik Indonesia [Kemenko PMK]. (2026a, March 12). *Surat Keputusan Bersama 7 Menteri Nomor 400.1-492 Tahun 2026 tentang Pedoman Pemanfaatan dan Pembelajaran Teknologi Digital dan Kecerdasan Artifisial di Jalur Pendidikan Formal, Nonformal, dan Informal*.

Kementerian Koordinator Bidang Pembangunan Manusia dan Kebudayaan Republik Indonesia [Kemenko PMK]. (2026b, March 13). *Pemerintah tetapkan SKB 7 Menteri untuk tata kelola AI di pendidikan*. https://www.kemenkopmk.go.id/pemerintah-tetapkan-skb-7-menteri-untuk-tata-kelola-ai-di-pendidikan-perkuat-koordinasi-dan

Kompas.com. (2026, March 13). *7 menteri sahkan aturan penggunaan AI dalam pendidikan*. https://tekno.kompas.com/read/2026/03/13/10280027/7-menteri-sahkan-aturan-penggunaan-ai-dalam-pendidikan

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). *Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task* [Preprint]. MIT Media Lab. https://www.media.mit.edu/publications/your-brain-on-chatgpt/

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–22). ACM. https://doi.org/10.1145/3706598.3713778

Li, Z., & Zhang, W. (2025). Technology in education: Addressing legal and governance challenges in the digital era. *Education and Information Technologies*, *30*(7), 8413–8443.

Luke, A. (2002). Beyond science and ideology critique: Developments in critical discourse analysis. *Annual Review of Applied Linguistics, 22*, 96–110. https://doi.org/10.1017/S0267190502000053

McDonald, N., Johri, A., Ali, A., & Collier, A. H. (2025). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. *Computers in Human Behavior: Artificial Humans, 3*, Article 100121. https://doi.org/10.1016/j.chbah.2025.100121

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*. Advance online publication. https://doi.org/10.1080/01425692.2025.2502808

Organisation for Economic Co-operation and Development [OECD]. (2026). *Digital education outlook 2026: Exploring the effective uses of generative AI in education*. OECD Publishing.

Selwyn, N. (2024). On the limits of artificial intelligence (AI) in education. *Nordisk tidsskrift for pedagogikk og kritikk, 10*, 3–14. https://doi.org/10.23865/ntpk.v10.6062

Simpson, N. (2025). Framing AI in higher education: A critical discourse analysis of inclusivity and power in university AI guidance documents. *Educational Linguistics, 4*(1). https://doi.org/10.1515/eduling-2024-0007

Sudding, M. F. J., Sakkir, G., Sudding, & Fitri, Q. (2026). Negotiating institutional values in the era of artificial intelligence: A critical discourse analysis of university AI guidelines. *ELT Worldwide, 13*(2), 394–405. https://doi.org/10.26858/eltww.v13i2.89230

United Nations Educational, Scientific and Cultural Organization [UNESCO]. (2023). *Guidance for generative AI in education and research*. UNESCO Publishing.

Williamson, B., Molnar, A., & Boninger, F. (2024). *Time to pause hasty adoption of AI applications in education*. National Education Policy Center.