
A Contextual Audit of Bangladesh's National AI Policy Draft 2026–2030

Farzana Islam
North South University
fi.farzanaislam@gmail.com

Tasmiah Tahsin
Mayeesha
University of North
Texas
tasmiahtahsinmayeesha
@my.unt.edu

With
Apart Research

Abstract

National AI policies in the Global South risk satisfying international benchmarks while failing to address locally salient concerns. This report audits Bangladesh's National AI Policy 2026–2030 against seven ethics principles derived from 32 stakeholder interviews conducted in prior work, AI4Bangladesh (ICTD 2024). Using pairwise qualitative comparison, we find the policy aligns well with principles for which international governance templates exist, such as individual rights, lawfulness, and accountability but shows substantive gaps on AI literacy as a deployment precondition and structurally independent oversight. A clause-by-clause comparison of successive drafts reveals that five independence mechanisms present in Draft v1.1 were removed in v2.0. We additionally identify five internal tensions in the policy, including underspecified provisions around centralized governance, the existing Bangla AI ecosystem, minority language inclusion, children's data, and private sector data access.

1. Introduction

National AI policies in the Global South face a structural challenge: governance frameworks are frequently adapted from regulatory templates developed in the Global North, without adequate grounding in local contexts, infrastructure realities, or stakeholder values [1, 2]. This creates policies that may satisfy international benchmarks while failing to address the specific risks most salient to the populations they govern.

Bangladesh released its National AI Policy 2026–2030 (Draft v2.0) in February 2026, following a multi-year process of stakeholder engagement and successive drafting. Separately, our prior work “AI4Bangladesh: AI Ethics for Bangladesh - Challenges, Risks, Principles, and Suggestions” conducted 32 semi-structured expert interviews with policymakers, ML practitioners, academics, and civil society stakeholders in Bangladesh, producing seven ethics principles grounded in local concerns [3]. These two bodies of work provide a natural opportunity for comparison: to what extent does the national policy reflect the priorities that emerged from structured stakeholder engagement with Bangladeshi experts?

This report presents a principled audit of Draft v2.0 against the seven AI4Bangladesh principles. Rather than scoring or ranking, we conduct a pairwise qualitative comparison — examining where the policy aligns with stakeholder-grounded principles, where it partially addresses them, and where meaningful gaps remain. We also compare the policy against the international frameworks it explicitly commits to following, using these not as external impositions but as the government's own stated benchmarks.

Our primary contributions are:

- A structured comparison of a national AI policy against locally grounded stakeholder research, offering a methodology applicable to other Global South contexts.
- Identification of areas of alignment, partial coverage, and absence between the AI4Bangladesh principles and Draft v2.0.
- Documentation of a structural change between successive drafts that has implications for the policy's commitments to independent oversight.
- Identification of five internal tensions and underspecified provisions within Draft v2.0 that warrant attention before finalization.

2. Related Work

The relationship between stakeholder-grounded AI ethics research and national policy development has received limited attention, particularly in Global South contexts. Prior work has examined AI ethics frameworks comparatively [4, 5], finding that transparency, fairness, and

accountability appear across most frameworks, while region-specific concerns — surveillance, data exploitation by foreign actors, and labor rights for data annotators — are underrepresented. Non-western AI ethics guidelines have been shown to share global values while also carrying locally distinct priorities that international frameworks do not capture [6].

In South Asia specifically, AI ethics research has emphasized the need for cultural re-contextualization of fairness research in language technologies [7], the particular risks of applying Global North fairness frameworks to local contexts [8], and the importance of explainability methods adapted to low-literacy populations [9]. Bangladesh's AI policy landscape has remained understudied; our earlier work [3] represents one of the first structured stakeholder studies of AI ethics priorities in the country.

Policy audit methodologies exist in related domains — algorithmic accountability [10], AI governance [11] — but structured comparisons between stakeholder-derived principles and national policy text remain uncommon. This work contributes to such a comparison for Bangladesh.

3. Methods

Source material. The audit draws on three documents: (1) AI4Bangladesh (ICTD 2024) [3], which produced seven ethics principles from 32 stakeholder interviews; (2) the Bangladesh National AI Policy 2026–2030, Draft v2.0 (February 9, 2026) [12]; and (3) Draft v1.1 (early February 2026), used to examine changes between successive versions of the policy. We also reference the four international frameworks the policy explicitly commits to in §4.10: the EU AI Act, the OECD AI Principles, the UNESCO Recommendation on AI Ethics, and the Council of Europe AI Framework Convention.

The seven AI4Bangladesh principles. These emerged from thematic analysis of 32 expert interviews conducted in 2023–2024 with stakeholders spanning computer science, law and policy, telecom, ML engineering, XAI research, and government administration. The principles are: (1) Individual Rights, (2) Lawfulness, (3) Fairness and Inclusivity, (4) Transparency and Explainability, (5) Accountability, (6) Literacy and Awareness, and (7) Oversight. Each was grounded in specific interview findings rather than derived from existing frameworks.

Comparison procedure. For each principle, we located the relevant sections of Draft v2.0, read them against the specific concerns that motivated the principle in the original stakeholder research, and characterized the relationship as one of: alignment (the policy addresses the concern with actionable provisions), partial coverage (the concern is named but lacks enforcement mechanisms or specificity), or absence (the concern is not addressed). We did not apply numeric scores. Where

the policy's text changed between v1.1 and v2.0 in ways relevant to a principle, we document this explicitly.

4. Results

4.1 Overview of Alignment

Table 1 summarizes the pairwise comparison across the seven principles. Three principles show strong alignment with the policy's provisions, two show partial coverage, and two show substantive gaps.

Table 1: Comparison of AI4Bangladesh Principles Against Draft v2.0

Principle	Core stakeholder concern	Policy coverage	Key clause(s)
Individual Rights	Data misuse, breach, privacy, harm to marginal communities	Aligned	§ 4.2, § 4.6, § 4.11
Lawfulness	No AI-specific legal framework; accountability gap	Aligned	§ 4.1, § 4.11, § 7.5
Fairness & Inclusivity	Bias in training data; exclusion of rural/low-literate users	Partial	§ 4.9, § 5.1
Transparency & Explainability	Opaque systems; civil society excluded from policy design	Partial	§ 4.6, § 3
Accountability	No enforcement culture; audit mechanisms absent	Aligned	§ 4.1, § 6.1
Literacy & Awareness	AI deployed before public can understand or contest it	Absent as precondition	§ 4.7
Oversight	Independent oversight essential given surveillance concerns	Structurally weakened	§ 6.1

4.2 Principle-by-Principle Comparison

Individual Rights. Stakeholder interviews raised concerns about data breaches (including a 2023 breach exposing 50+ million citizens' records), foreign companies extracting Bengali speech data without compensation, and AI systems that could harm marginalized populations. Draft v2.0 addresses these concerns substantively. The Personal Data Protection Ordinance (PDPO) 2025 is incorporated by reference; § 4.2 introduces data localization requirements and algorithmic disgorgement for unlawfully obtained training data; § 4.6 establishes rights-based protections against automated decisions, including the right to counterfactual explanation and human review. § 4.11 introduces strict liability for high-risk AI systems. These provisions map closely to the stakeholder-identified concerns.

Lawfulness. A consistent finding across interviews was the near-complete absence of AI-specific law in Bangladesh, leaving harms legally unaddressed under existing ICT frameworks. Draft v2.0 takes a layered approach: existing ordinances (PDPO, National Data Governance Ordinance, Cyber Safety Ordinance 2025) provide an immediate foundation, while § 7.5 mandates development of comprehensive AI liability legislation by 2028. A risk-based classification framework (§ 4.1) establishes categories of prohibited, high-risk, limited-risk, and low-risk systems. The policy aligns with this principle, though the 2028 legislative timeline means enforcement gaps will persist for several years.

Fairness and Inclusivity. Stakeholders emphasized the importance of inclusive training data, protection of marginalized communities, and designing AI products for low-literacy users. Draft v2.0 names exclusion risks explicitly in § 4.9, listing women, elderly, economically disadvantaged, religious and linguistic minorities, and disabled people as groups whose exclusion must be "thoroughly examined and prevented." § 5.1 requires training datasets to be "inclusive and representative." However, these provisions are largely aspirational; no binding rural equity safeguard, specific threshold, or non-digital fallback mechanism is defined. The gap between naming a population and establishing an enforceable obligation to serve them represents a meaningful limitation, particularly given that 58% of Bangladesh's population lives in rural areas with significantly lower internet access.

Transparency and Explainability. Stakeholder concerns spanned two distinct levels: transparency of AI systems (explainability to affected users) and transparency of the policy design process itself. Draft v2.0 addresses the former reasonably well § 4.6 requires counterfactual explanations for high-stakes automated decisions, framed in terms of outcomes rather than technical disclosures. § 3 (AI Principles) includes "transparent and responsible" as a stated value. However, the second concern that AI strategy in Bangladesh had historically been developed without meaningful civil society or academic input is addressed only weakly. Public consultations are referenced in the 2028 mid-term review (§ 7.2) but are not built into deployment decisions or annual reporting as a

mandatory recurring mechanism. Our stakeholders specifically requested that policy design transparency be treated as distinct from system-level transparency; the policy collapses them.

Accountability. Interviews found that accountability in Bangladesh was effectively user-centric — companies faced minimal institutional liability. Draft v2.0 represents a substantive shift: § 4.1 mandates Algorithmic Impact Assessments (AIAs) for public sector and high-risk private AI systems; § 6.1 establishes an oversight committee with authority to audit deployments for bias, error rates, and red-teaming results; § 4.11 introduces strict liability for high-risk systems with a five-year evidence preservation requirement. These provisions align with what stakeholders called for, representing meaningful progress on institutional accountability.

Literacy and Awareness. This principle was distinctive to AI4Bangladesh; it does not appear prominently in most international frameworks. Stakeholders were concerned that AI systems were being deployed into a population where AI literacy was extremely low, leaving citizens unable to recognize when automated systems were affecting them, let alone contest such decisions. Draft v2.0 includes public awareness campaigns (§ 4.7) and civil service capacity building (§ 4.8), but treats literacy as a downstream communication effort rather than a precondition for deployment. The policy does not establish any literacy or notification threshold that must be met before high-stakes AI systems are rolled out in a given area. The principle's intent — that a population must be able to understand and contest an AI decision before such a system is deployed at scale — is not operationalized.

Oversight. This principle carried particular weight in the stakeholder research. Many participants had lived through the 2024 Quota-protest period, during which authorities used internet shutdowns, surveillance, and monitoring to identify protesters. The concern was not simply that AI would be misused, but that a government-controlled oversight body would be structurally unable to check government misuse. Participants repeatedly called for independent oversight bodies at arm's length from both government and industry.

Draft v2.0 § 6.1 opens with language that directly echoes this: the committee "shall operate at arm's length from both government and industry." However, a comparison with Draft v1.1 reveals that between the two versions — separated by days in early February 2026 — five structural mechanisms of independence were removed. Table 2 documents these changes.

Table 2: Changes to § 6.1 (Independent Oversight Committee) Between Drafts v1.1 and v2.0

Mechanism	Draft v1.1	Draft v2.0
Legislative basis	Created by Act of Parliament	Not specified
Budget autonomy	Parliamentary budget independence	Not specified
Tenure and nominations	Fixed tenure; public call for nominations	Not specified
Conflict-of-interest	Disclosure required	Not specified
Institutional location	Coordinates with NDGA; sits outside it	Formed within the Executive Council of the NDGA

The phrase "arm's length" was retained; the provisions that would operationalize it were not. Additionally, Draft v2.0 § 4.10 commits Bangladesh to ratifying the Council of Europe AI Framework Convention, which requires a legally independent oversight mechanism. The current § 6.1 structure, as drafted, does not satisfy this requirement.

4.3 Additional Policy Gaps: Internal Tensions and Underspecified Provisions

Beyond the principle-level comparison, a close reading of Draft v2.0 reveals several internal tensions and underspecified provisions that raise questions about implementation feasibility. These are distinct from gaps identified against the AI4Bangladesh principles; they are issues internal to the policy's own logic.

Centralization assumptions and institutional stability. A structuring premise of Draft v2.0 is that a centralized governmental body—primarily the National Data Governance Authority (NDGA)—will coordinate AI policy across government, academia, and industry while maintaining ethical standards, data protection, and cybersecurity frameworks. This premise deserves scrutiny in Bangladesh's current institutional context. The student-led July 2024 revolution that the policy itself references in § 1.1 as a catalyst for reform also demonstrated that government institutions in Bangladesh can undergo rapid and significant change. A governance architecture that concentrates coordination authority within a single government-controlled body carries meaningful continuity risk when the political environment is unsettled. The policy does not address several practical questions that follow from this design: what happens to ongoing AI governance work if the NDGA undergoes restructuring; whether employees of joint government-industry AI initiatives would hold civil service or private sector status; and what accountability mechanisms apply to the coordinating body itself when its institutional position changes.

Omission of the existing AI ecosystem. Draft v2.0 proposes the development of a national Bangla language model as a government-led initiative (§ 6.3), framing it as a foundational project. This proposal does not acknowledge that a substantial ecosystem of Bangla language models already exists, developed by university researchers, domestic technology companies, and diaspora-connected teams over the preceding decade. Similarly, the policy does not reference existing government AI initiatives — including those developed under the a2i (aspire to innovate) program and other ICT Division projects — despite these representing years of institutional investment and accumulated technical knowledge. A national AI policy that does not map or build on the existing landscape risks duplicating work, creating competing standards, and signaling to established researchers and developers that their contributions are not recognized. The policy's silence on the existing ecosystem is a missed opportunity to anchor new initiatives in what has already been learned.

Linguistic and cultural diversity. The policy correctly identifies linguistic diversity as central to Bangladesh's AI future (§ 1.3) and mandates that AI systems support Bangla and other "nationally relevant languages." However, it does not name or make specific provision for the languages of Bangladesh's ethnic and indigenous minority communities — including the languages spoken in the Chittagong Hill Tracts and other regions with distinct linguistic identities. Similarly, while the policy invokes cultural context as a value, it does not specify how cultural diversity beyond the Bengali majority context should be reflected in AI system design, training data, or evaluation. The gap between the principle of cultural inclusivity and its operationalization is particularly significant for AI systems in public services, where minority-language communities are already underserved.

Children's data and the safety-utility tension. § 4.1 prohibits the collection and use of children's data (under 12) for commercial purposes and for training AI models. The prohibition on commercial use is straightforward. The prohibition on using children's data for AI training, however, does not account for cases where such data is necessary to develop AI systems that protect children — for example, child safety classifiers, content moderation systems, pediatric health diagnostics, or educational tools designed specifically for young learners. Effective AI safety and child protection tools often require training on representative data from the populations they are designed to serve. The policy draws a categorical line without distinguishing between exploitative

commercial use and safety-motivated or public-interest training, creating a provision that may inadvertently impede the development of child-protective AI systems. A more nuanced framework would distinguish by purpose, oversight mechanism, and consent regime rather than applying a blanket prohibition on data use for model training.

Private data in the AI Data Exchange. § 4.2 establishes a government-managed AI Data Exchange under the NDGA to facilitate access to high-quality datasets for research, innovation, and public-interest applications. The provision describes privacy-preserving mechanisms (anonymization, aggregation, consent management) and defines access controls for researchers, startups, academia, and government agencies. What the provision does not address is the status of data held by private companies—including telecommunications providers, ride-sharing platforms, financial services companies, and e-commerce operators—which collectively hold some of the most behaviorally rich and demographically representative datasets in Bangladesh. The policy's data governance framework is structured around public sector data and cross-border transfer of sensitive categories; it does not establish a mechanism through which privately held data could be made available for public-interest AI research under appropriate safeguards. This omission limits the practical utility of the AI Data Exchange, since the most valuable training datasets for many socially relevant applications will not be in government hands.

5. Discussion and Limitations

5.1 What the Comparison Shows

The audit reveals a pattern that is not uniform: Draft v2.0 performs well on principles that map onto existing international governance machinery. Individual rights protections, liability frameworks, and accountability mechanisms have clear templates in the EU AI Act and OECD Principles, and the policy adapts these competently. The gaps appear precisely where Bangladesh's own context provides no international template - AI literacy as a deployment precondition, non-digital fallbacks for a majority-rural population, and oversight structures independent of a government with a recent history of surveillance.

This pattern has implications beyond Bangladesh. It suggests that Global South AI policies risk systematically underproviding on the dimensions most relevant to their populations, not because of neglect, but because the available governance templates do not encode those dimensions. Locally grounded stakeholder research is therefore not supplementary to policy development, rather it is the primary source for identifying what international frameworks leave out.

The oversight regression between v1.1 and v2.0 is worth noting precisely because it is not a case of an absent provision, the language of independence was present and then structurally hollowed out. This is distinct from a gap; it represents a redrafting choice whose implications warrant attention before the policy is finalized.

5.2 Limitations

This audit is based on the February 9, 2026 version of Draft v2.0; subsequent revisions may alter the findings. The comparison is qualitative and reflects interpretive judgment about whether policy text meaningfully addresses stakeholder concerns, but reasonable analysts might characterize some provisions differently. The AI4Bangladesh principles were derived from 32 interviews, which, while sampled for breadth, cannot claim full representativeness of Bangladesh's population. Finally, this audit addresses policy text, not implementation, a well-written provision may not translate to practice, and a partially specified provision may be operationalized through secondary regulations not yet published.

5.3 Future Work

Future work could apply the same approach to other South Asian or Global South AI policies as they emerge, building a comparative picture of where locally grounded concerns do and do not surface in national frameworks. Within Bangladesh, a follow-up audit of implementation against the finalized policy would be a natural next step.

6. Conclusion

We conducted a principled pairwise comparison between seven stakeholder-grounded AI ethics principles from AI4Bangladesh (ICTD 2024) [3] and Bangladesh's National AI Policy 2026–2030 (Draft v2.0) [12]. The policy aligns well with principles for which international governance templates exist - individual rights, lawfulness, and accountability - and shows partial coverage on fairness and transparency. Two principles central to Bangladesh's specific context remain substantively unaddressed: AI literacy as a precondition for deployment, and structurally independent oversight. A comparison of successive drafts shows that oversight independence was present in v1.1 and removed in v2.0, while the commitment to the Council of Europe AI Convention - which requires independent oversight - was retained.

Beyond the principle comparison, we identify five internal tensions in the policy: an assumption of centralized institutional stability that may not hold in Bangladesh's current context; omission of the existing Bangla AI ecosystem; a gap between stated commitments to linguistic and cultural diversity and their operationalization for minority communities; a categorical prohibition on children's data for AI training that does not account for child-protective use cases; and an AI Data Exchange framework that does not address privately held data. These gaps are not failures of

intent; the policy is substantively ambitious but they represent provisions that, if left unresolved, may limit the policy's practical effectiveness or produce unintended consequences. The audit method itself, benchmarking national policy against locally derived principles and the government's own stated commitments, offers a reusable approach for researchers in comparable contexts.

Code and Data

This project involves policy analysis; no code repository applies.

References

- [1] Arun, C. 2019. AI and the global south: Designing for other worlds. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES 2019).
- [2] Png, M.-T. 2022. At the tensions of south and north: Critical roles of global south stakeholders in AI governance. In Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT 2022), 1434–1445.
- [3] Mayeesha, T.T., Islam, F., and Ahmed, N. 2024. AI4Bangladesh: AI Ethics for Bangladesh Challenges, Risks, Principles, and Suggestions. In Proceedings of the 13th International Conference on Information & Communication Technologies and Development (ICTD 2024). <https://doi.org/10.1145/3700794.3700820>
- [4] Jobin, A., Ienca, M., and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1(9), 389–399.
- [5] Hagendorff, T. 2020. The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines* 30(1), 99–120.
- [6] Hongladarom, S. and Bandasak, J. 2023. Non-western AI ethics guidelines: implications for intercultural ethics of technology. *AI & Society* (2023), 1–14.
- [7] Bhatt, S., Dev, S., Talukdar, P., Dave, S., and Prabhakaran, V. 2022. Cultural Re-contextualization of Fairness Research in Language Technologies in India. In Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the ACL and 12th International Joint Conference on Natural Language Processing (AACL-IJCNLP 2022), 727–740.
- [8] Sambasivan, N., Arnesen, E., Hutchinson, B., Doshi, T., and Prabhakaran, V. 2021. Re-imagining algorithmic fairness in india and beyond. In Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT 2021), 315–328.

- [9] Okolo, C.T., Dell, N., and Vashistha, A. 2022. Making AI explainable in the Global South: A systematic review. In Proceedings of the ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS 2022), 439–452.
- [10] Raji, I.D., Smart, A., White, R.N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., and Barnes, P. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT 2020), 33–44.
- [11] Schiff, D., Biddle, J., Borenstein, J., and Laas, K. 2020. What's next for AI ethics, policy, and governance? A global overview. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES 2020), 153–158.
- [12] Government of Bangladesh, ICT Division. 2026. National Artificial Intelligence Policy 2026–2030, Draft v2.0. February 9, 2026. Available at: <https://aipolicy.gov.bd/>

LLM Usage Statement

Claude was used in improving writing and formatting. The analysis was done by the authors.