

A National AI-in-Education Policy: What Malaysia Should Write Before Vendors Write It for Us

Hafiz Hanif

Faculty of Human Development, Sultan Idris Education University · ORCID 0000-0001-9943-6557

CITE AS	Hanif, H. (2026, September 5). A National AI-in-Education Policy: What Malaysia Should Write Before Vendors Write It for Us. Hack(edu).tech Chronicle. https://hackedu.tech/articles/a-national-ai-in-education-policy-what-malaysia-should-write-before-vendors-writ
PUBLISHED	5 September 2026
DOI	10.5281/zenodo.22331347
LICENCE	CC-BY-4.0
AVAILABLE AT	https://hackedu.tech/articles/a-national-ai-in-education-policy-what-malaysia-should-write-before-vendors-writ

ABSTRACT

As AI vendors court Malaysian schools, this article asks what policy instrument governs their entry. Crediting the real groundwork laid, MOSTI's AIGE principles, the Kementerian Digital's public-sector AI adoption guideline with its four risk tiers, KPT's generative AI guideline for higher education, and KPM's 2026 Panduan Literasi Kecerdasan Buatan, it argues that a panduan is nonetheless not a dasar: guidance advises where policy binds. The missing instrument must govern procurement, vendor accountability, children's data, assessment redesign, and funded teacher training. The article proposes a ten-section draft national AI-in-education policy and calls for it to be debated publicly, before vendors write it by contract.

KEYWORDS

Policy, Artificial Intelligence

The vendors are already in our Bilik Guru.

I have sat in enough demo sessions to know the script by heart. The slides are beautiful. The AI tutor speaks politely. The dashboard glows with numbers. Somewhere near the end comes the promise that this platform will "transform" our schools, and somewhere after that comes the price, usually per student, per year, forever.

I do not blame the vendors. Selling is their job. My question is a different one: when they arrive at the ministry's door, at the JPN, at the PIBG meeting, what document do we hold in our hands to judge them? What is our policy on AI in education?

Right now, the honest answer is: we have pieces, but we do not have the policy. And in policy, a vacuum never stays empty. Someone will fill it. If we do not write our own rules, the rules will arrive pre-written, in the terms and conditions of whichever platform signs the contract first.

Let me be fair first, because credit must be given where it is due, and more has been done than many people realise.

The Existing Documents

At the national level, Malaysia is not sleeping. We have the National AI Roadmap 2021-2025, and in September 2024 MOSTI launched the National Guidelines on AI Governance and Ethics, known as AIGE, which sets out seven principles for responsible AI, covering fairness, reliability, privacy, and accountability, and names education as one of its key sectors (MOSTI 2024). A National AI Office has been set up. In March 2025, the Kementerian Digital issued its guideline for AI adoption across the public sector, and it is a more serious document than its quiet circulation suggests: it defines the roles of every actor from developer to end user, sets adoption requirements covering data quality, training, and budget planning, provides a self-assessment instrument for ethics compliance, and classifies AI applications into four risk tiers, from prohibited to minimal risk, in the spirit of the European Union's AI Act (Kementerian Digital 2025). These are real foundations.

The Ministry of Higher Education has gone further, and I want to acknowledge this properly. KPT, through Jabatan Pendidikan Tinggi, has issued the Garis Panduan Penggunaan Teknologi Kecerdasan Buatan Generatif dalam Pengajaran dan Pembelajaran for our universities (Kementerian Pengajian Tinggi (KPT) 2024). Whatever its limitations, it exists. It tells lecturers and students that generative AI is a reality to be managed, not a sin to be denied. Our IPTs at least have a reference point when a student asks:

"boleh ke saya guna ChatGPT untuk assignment ni?"

KPM, the ministry responsible for our schools, has also been moving, and having now read its newest document from cover to cover, I want to give it proper credit. In 2026, the ministry published the Panduan Literasi Kecerdasan Buatan (AI), developed by Bahagian Sumber dan Teknologi Pendidikan and given administrative force through Surat Pekeliling Ikhtisas Bil. 2 Tahun 2026 (Kementerian Pendidikan Malaysia (KPM) 2026). This is not a thin pamphlet. Its framework spans four domains, AI governance, teaching and learning, lifelong learning, and AI ethics. It defines roles from school leaders to parents to national agencies. It touches assessment adaptation, references the PDPA and the Education Act, includes a phased action plan with monitoring, and even provides guidance for selecting AI tools. It sits alongside the Digital Education Policy of 2023, the AI-Powered Classroom pilots, and AI basics entering the curriculum. Taken together, this is the most serious groundwork our school system has laid so far, and the team behind it deserves to hear that.

So my argument today is not that KPM has done nothing. It is that what exists, good as it is, is still a 'panduan', and a 'panduan' is not a 'dasar'. The difference is not academic; it is the difference between advice and obligation. A panduan tells teachers, leaders, and officers how to behave with AI, and this one does that well. A policy tells the system how it must behave: which products may enter, on what evidence, under what contracts, with whose money, audited by whom, with what consequences when the rules are broken. Read the new Panduan carefully and you will see it points at these matters, urging compliance with existing laws, advising careful tool selection, but pointing is not binding. Where the Panduan says choose tools wisely, a policy says no tool enters our schools without meeting these published criteria. Where the Panduan reminds us the PDPA exists, a policy writes children-specific data rules with an audit mechanism behind them. That binding instrument, covering procurement, vendor obligations, data enforcement, assessment redesign at national level, and funded teacher training in one document, is the book that still needs writing, and it should be written in the open, because a policy for five million students should not be composed quietly.

The Gap

Why does the gap matter so much? Three reasons, and all three we have paid for before.

- First, procurement without policy produces shelfware. We have lived this story. Devices bought without asking who maintains them. Platforms licensed without asking whether teachers were trained. Our own history of grand educational technology initiatives that faded quietly should have taught us that buying is the easiest part and the least important part. AI raises the stakes because now the product is not a dusty tablet; it is a system that talks to our children and learns from their data. Research on AI for teachers keeps finding the same thing: the technology's promises run ahead of what teachers actually need in classrooms (Celik et al. 2022). Without a policy that forces the question "what problem does this solve, and what is the evidence?", we will buy the demo, not the outcome.
- Second, the data. Every AI platform that enters a school becomes a collector of children's information: their mistakes, their pace, their writing, sometimes their voices and faces. Reviews of large language models in education list privacy and data governance among the most serious unresolved challenges (Yan et al. 2023). Our PDPA was not written with school-going minors and foreign AI platforms in mind. Who owns the learning data of a Standard Three pupil? Where is it stored? Can the vendor train future models on it? Can they sell insights derived from it? Today, the answer sits in each vendor's terms of service, written by their lawyers, for their protection. That is what I mean when I say the policy is being written for us.
- Third, the pedagogy and the pushback. Here I want to speak honestly about something I have observed in the past year or two, including among fellow academics. The early excitement about AI in learning has met a real resistance, and it is not ignorant resistance. It is grounded in evidence that heavy, unguided reliance on AI tools is associated with weaker critical thinking through cognitive offloading, especially among younger users (Gerlich 2025). Lecturers are redesigning assessments because they can no longer tell which submissions were thought and which were generated. Teachers worry, reasonably, that the struggle that produces learning is being quietly outsourced. This pushback deserves respect, not dismissal. But fear without policy produces the worst of both worlds: blanket bans in one school, free-for-all in the next, and students everywhere learning that the adults have no coherent answer. The research is actually quite clear on the direction: reviews of AI-powered educational agents find that hybrid arrangements, where teachers stay in the loop curating and moderating, outperform fully autonomous systems (Córdova-Esparza 2025). Policy is how we turn that finding into practice at national scale.

And there are questions beyond these three that only a national policy can answer. Teacher readiness: studies of our own region show that meaningful technology integration depends on teachers' blended knowledge of content, pedagogy, and technology, which does not appear by circular; it is built by sustained training (Sofwan et al. 2023). Who trains 400,000 teachers, on what standard, funded how? Classroom support: when the AI platform fails during period three, who does the teacher call? Language: as I wrote in an earlier column, the machinery itself taxes Bahasa Melayu, performing worse and costing more in our language than in English (Petrov et al. 2023), so a policy that ignores language will quietly favour English-medium learning by default. Equity: the B40 student and the international school student must not end up with two different qualities of AI. And underneath everything, the foundational research on AI literacy tells us the ordering: understanding what AI is must come before using it (Ng et al. 2021), which means the new Panduan, with literacy at its core, is exactly the right first book. My argument is simply that it cannot be the only one.

International guidance already exists to help us. UNESCO has published guidance for generative AI in education, urging governments to regulate before scaling, to protect data, to set age and validation standards, and to keep human agency at the centre (UNESCO 2023). The large research reviews spell out both the opportunities and the risks in detail (Kasneci et al. 2023). We do not need to invent from zero. We need to decide, in our own context, with our own values, in our own language.

The Proposed Policy

So let me do what columnists usually avoid, and put a concrete draft on the table. Not because I think a lecturer can write a national policy alone, but because a debate needs a starting document to argue with. Here is my proposed skeleton for a Dasar Kecerdasan Buatan dalam Pendidikan, a National AI-in-Education Policy, in plain point form.

1. Purpose and principles

- State clearly why AI is allowed in our schools: to reduce teacher burden, to widen access to quality learning support, and to prepare students for an AI-shaped economy and society.
- Anchor the policy to existing foundations: the Falsafah Pendidikan Kebangsaan, the AIGE principles, and the constitutional position of Bahasa Melayu.
- Declare the non-negotiables: the teacher remains the pedagogical authority; AI supports, never replaces, human judgment on any decision that affects a child's future.

2. AI literacy and competency standards

- Adopt the KPM Panduan Literasi Kecerdasan Buatan (AI) and its four-domain framework as the literacy baseline for all students and school leaders, staged by school level.
- Define a teacher competency framework: what every teacher must know (how AI works, its limits, verification habits), what specialist teachers must know, and what school leaders must know for procurement decisions.
- Make AI literacy a compulsory component in teacher education institutes and in-service training, with protected time, not another burden squeezed into weekends.

3. Teaching and learning use

- Define permitted uses by category: preparation and planning, materials generation, feedback support, translation and simplification, administrative work.
- Require disclosure norms: teachers and students state when and how AI was used, normalising honesty instead of policing suspicion.
- Require that AI-generated teaching materials remain subject to teacher review before classroom use, consistent with evidence favouring teacher-in-the-loop models.

4. Assessment integrity and redesign

- Direct the examination and curriculum bodies to review national and school-based assessment formats against AI capability, prioritising process evidence, oral defence, in-class performance, and locally-anchored tasks.
- Prohibit reliance on AI-detection tools as sole evidence in misconduct cases, given their known unreliability.
- Fund and publish Malaysian research on AI-resilient assessment so the redesign is evidence-based, not panic-based.

5. Procurement and vendor accountability

- Build on the Kementerian Digital's 2025 guideline for AI adoption in the public sector, extending its principles into binding, education-specific procurement rules rather than starting from a blank page.
- Adapt its four risk tiers to schooling and attach obligations to each: prohibited uses (emotion surveillance of pupils, indiscriminate collection of children's biometric data, manipulative engagement design); high-risk uses (AI involved in grading national assessments, streaming, or disciplinary decisions) permitted only with mandatory human review, published validation, and audit; limited-risk uses (chatbots and tutoring tools) requiring disclosure and age-appropriate safeguards; minimal-risk uses (drafting worksheets, formatting) requiring only baseline compliance.
- No AI product enters schools without meeting published criteria: demonstrated learning evidence, Bahasa Melayu performance benchmarks, data practices compliant with this policy, offline or low-bandwidth degradation plans, and local support commitments.

- Require pilot evaluation with independent measurement before any statewide or national rollout, and require the results to be published.
 - Ban contract terms that allow vendor training on student data without explicit, revocable, ministry-level consent.
 - Set a sunset rule: every licensed platform is re-evaluated on evidence every two or three years, with a clear exit clause.
6. Student data protection
- Classify school learners' data as protected children's data with stricter rules than general PDPA requirements.
 - Require data residency or equivalent legal safeguards, plain-language parental notification, and the right to deletion upon leaving school.
 - Establish an audit mechanism: the ministry, not the vendor, verifies compliance.
7. Language and equity
- Require Bahasa Melayu capability testing for any AI tool deployed in national schools, with published benchmark results.
 - Support and give procurement weight to local and regional models and content, so the ecosystem grows.
 - Guarantee baseline access: AI-supported learning provisions must reach B40 students, rural schools, and students with special needs on an equal footing, with connectivity and device gaps addressed in the same policy, not postponed to another one.
8. Infrastructure, support, and the teacher's working day
- Every deployment plan must name the support model: who maintains, who troubleshoots, who trains, at what ratio, funded from which line.
 - Apply a subtraction rule, one I have argued for many times: for every new AI system introduced, an existing administrative requirement is retired, and the ministry reports annually on net hours returned to teachers.
9. Governance, monitoring, and review
- Establish a standing AI-in-Education council with teachers, parents, academics, ministry officials, and student representation, not vendors, as voting members.
 - Publish an annual public report: what was deployed, what was measured, what was stopped.
 - Review the whole policy every three years, because this technology will not wait for our ten-year planning cycles.
10. Research and evidence pipeline
- Fund Malaysian research on AI in our classrooms, in our languages, with our curriculum, through existing grant structures, and require deployments to cooperate with independent researchers.
 - Partner IPTs as regional evaluation and training hubs, connecting the expertise we already have to the schools that need it.

That is my draft. It is surely incomplete, possibly wrong in places, and absolutely open for others to tear apart. That is the point. A policy of this weight should be argued about in public, in the newspapers, in the Dewan, in staffrooms and PIBG meetings, not composed quietly and discovered later in a procurement announcement.

I write this with genuine respect for the people in KPM who produced the Panduan; having read it, I can say the foundation is better than many of us assumed. But a foundation is laid so that something can be built on it, and our children need the whole building, because the vendors are not waiting for us to finish it.

We like to say that education is our nation's future. If we believe it, then the rules for the most powerful educational technology of our lifetime cannot be left as an appendix to someone's sales contract.

Write the policy. Write it openly. Write it now, while the ink is still ours to hold.

References

- Celik, I., Dindar, M., Muukkonen, H., & Järvelä, S. (2022). The Promises and Challenges of Artificial Intelligence for Teachers: a Systematic Review of Research. *TechTrends*, 66(4), 616–630. <https://doi.org/10.1007/s11528-022-00715-y>
- Córdova-Esparza, D.-M. (2025). AI-Powered Educational Agents: Opportunities, Innovations, and Ethical Challenges. *Information*, 16(6), 469. <https://doi.org/10.3390/info16060469>
- Gerlich, M. (2025). AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kementerian Digital. (2025). Garis panduan pengadaptasian kecerdasan buatan (AI) sektor awam (Edaran 4 Mac 2025). Kementerian Digital Malaysia.
- Kementerian Pendidikan Malaysia (KPM). (2026). Panduan literasi kecerdasan buatan (AI) Kementerian Pendidikan (Surat Pekeliling Ikhtisas KPM Bil. 2 Tahun 2026). Kementerian Pendidikan Malaysia. <https://www.moe.gov.my/storage/files/shares/Pengumuman/BUKU%20PANDUAN%20LITERASI%20KECERDASAN%20BUATAN%20AI%209.6.26.pdf>
- Kementerian Pengajian Tinggi (KPT). (2024). Garis panduan penggunaan teknologi kecerdasan buatan generatif (KBG) dalam pengajaran dan pembelajaran (PdP) pendidikan tinggi. Jabatan Pendidikan Tinggi.
- MOSTI. (2024). National guidelines on AI governance and ethics (AIGE). Ministry of Science, Technology and Innovation Malaysia.
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- Petrov, A., La Malfa, E., Torr, P. H. S., & Bibi, A. (2023). Language model tokenizers introduce unfairness between languages. *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, 36963–36990.
- Sofwan, M., Habibi, A., & Yaakob, M. F. M. (2023). TPACK's Roles in Predicting Technology Integration during Teaching Practicum: Structural Equation Modeling. *Education Sciences*, 13(5), 448. <https://doi.org/10.3390/educsci13050448>
- UNESCO. (2023). Guidance for generative AI in education and research. United Nations Educational, Scientific and Cultural Organization.
- Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gašević, D. (2023). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112. <https://doi.org/10.1111/bjet.13370>