PROCEEDINGS THE 5th INTERNATIONAL CONFERENCE ON ECONOMICS, BUSINESS, AND MANAGEMENT RESEARCH (ICEBMR)
"Digital Economy, Human Dignity, and Inclusive Transformation in Asia"
https://econference.usd.ac.id/index.php/ICEBMR | ISSN: 3032-596X | Vol 5, 2026

AI FOR SOUTHEAST ASIA: CAPABILITY, CONTROL, AND RESPONSIBILITY FOR HUMAN FLOURISHING

William Chandra Tjhi[1], Antonius Sumarwan[2]
[1]AI Practitioner, Singapore
[2]Sanata Dharma University, Indonesia
[1]williamtjhi@gmail.com, [2]marwansj@usd.ac.id
DOI: https://doi.org/10.24071/icebmr.v5i1.821

Abstract
Artificial intelligence is entering everyday life in Southeast Asia through education, work, commerce, healthcare, civic participation, and disaster response. This article develops a regional perspective drawn from the author's experience of building open regional Large Language Models (LLMs) over the past few years, arguing that AI can narrow divides only when capability, control, and responsibility move closer to the communities most affected. This article organizes the discussion around four structural gaps - language and voice, cultural fit, infrastructure, and human support - and three risk domains: education, work, and trust. The analysis highlights the conditional nature of AI as an equalizer: lower-cost access to expertise does not automatically create agency. Edge AI, local-language resources, open evaluation, human oversight, and community participation therefore become design and governance considerations rather than optional additions. The article further cites SEA-LION (https://sea-lion.ai/blog/sea-lion-v4-multimodal/) as an ecosystem example of models, data, evaluation, experimentation, and community infrastructure, and develops a proposal for an ASEAN AI-for-Good institution organized around sustain, scale, and steward functions. It concludes that a positive AI future in Southeast Asia depends on collaboration and democratization of capability and meaningful control, with human dignity and local decision-making retained as practical tests.
Keywords: *artificial intelligence, Southeast Asia, edge AI, human agency, AI governance*

Introduction

Artificial intelligence is no longer experienced only as a frontier technology. In Southeast Asia, it is entering ordinary settings in which people learn, teach, trade, care, participate in civic life, and respond to disasters. The keynote that motivates this article frames the central question in practical rather than purely technical terms: AI can narrow the region's divides only when capability, control, and responsibility move closer to the communities most affected (Tjhi, 2026). This framing shifts attention from whether AI exists to where the ability to use, correct, govern, and sustain AI is located.

The regional setting makes this question especially consequential. Southeast Asia contains substantial linguistic and cultural diversity, uneven digital infrastructure, and different institutional capacities. The result is that the same AI technology can encounter very different constraints depending on language, connectivity, professional support, and the consequences of error. The presentation therefore treats the AI "equaliser" as conditional. Its potential includes on-demand translation, personalised learning, decision support, and lower-cost creation, while its frictions include connectivity and device constraints, language and voice gaps, limited skills and supervision, and trust and safety concerns (Tjhi, 2026).

The issue is not simply one of access. A person may technically reach an AI system without having meaningful agency over the system or the decisions that follow from its output. This article proposes four practical design questions: who can use the system, who can correct it, who bears the error, and who retains capability after the interaction (Tjhi, 2026). These questions connect technical design with organizational responsibility and, ultimately, with the distribution of decision-making power.

This article examines everyday points of AI entry in the region, identifies the structural constraints visible at the edge of deployment, discusses education, work, and trust as three domain examples in which risks acquire local forms, and uses the SEA-LION journey (https://sea-lion.ai/blog/sea-lion-v4-multimodal/) as an example of open regional capability building. It then develops the presentation's proposal for an ASEAN AI-for-Good institution and concludes with a set of

practical questions for assessing whether AI strengthens human dignity, expands local capability, preserves human judgment, and can be sustained collectively (Tjhi, 2026).

Literature Review

AI, Human Agency, and Responsible Governance

A useful way to understand human capability, control, and responsibility on the development and use of AI is to place it alongside existing regional and international approaches to AI governance. The ASEAN Guide on AI Governance and Ethics presents practical guidance for the responsible design, development, and deployment of AI in the region and emphasizes alignment and interoperability across jurisdictions (ASEAN Secretariat, 2024). Its guidance also identifies risks associated with inaccurate outputs, disinformation, deepfakes, impersonation, phishing, privacy, and confidentiality in generative AI contexts (ASEAN Secretariat, 2025). These themes are closely related to the keynote's concern that access to AI does not by itself produce safe or meaningful agency.

At the global level, UNESCO's Recommendation on the Ethics of Artificial Intelligence places human dignity and human rights at the centre of AI ethics and highlights transparency, fairness, diversity, inclusiveness, privacy, and human oversight (UNESCO, 2021). The alignment is important for the article's conceptual framing: human agency is not treated as a purely individual ability to operate a tool, but as a condition supported by social, institutional, and technical arrangements. In the Southeast Asian context presented by Tjhi (2026), this means that the people closest to a problem need not merely access a model; they also need the ability to evaluate, correct, contextualise, and govern its use.

Localisation, Language, and Cultural Fit

Language localisation is a central component of the regional argument. Language support is separated into text, voice, and culture. For text, it points to tokenizer inefficiency, code-switching, local terminology, and the need for robust evaluation. For voice, it highlights accents and dialects, noise and low bandwidth, speech-data consent, and on-device constraints. For culture, it identifies local institutions, norms and sensitivities, neighbour-language leakage, and community values as dimensions of understanding (Tjhi, 2026). This implies that multilingual performance cannot be reduced to whether a system can generate grammatically fluent sentences.

The SEA-LION ecosystem provides a concrete regional response to this localisation challenge. AI Singapore describes SEA-LION as an open model family developed around Southeast Asian linguistic and cultural needs, while SEA-HELM is designed to evaluate linguistic and cultural competencies in Southeast Asian languages (AI Singapore, 2025a, 2025b). The ecosystem also includes open language-data infrastructure, including Project SEALD and Project ATLAS, which are intended to improve access to high-quality regional language data and enable community contribution (AI Singapore, 2025c). Together, these efforts support the presentation's broader proposition that local capability must exist across data, models, evaluation, deployment, and especially people rather than being concentrated in a single model.

AI as a Conditional Equalizer

AI is presented as a possible equalizer because it can reduce the cost of accessing expertise. That potential is visible in educational support, business recommendations, translation, decision support, and content creation. However, these benefits are countered by some frictions such as device access, connectivity, language coverage, skills, supervision, and trust (Tjhi, 2026). This conditional formulation is consistent with the regional emphasis of AI Ready ASEAN, a programme designed to build AI literacy among youth, educators, parents, and underserved communities through localised learning and regional capacity building (ASEAN Foundation, 2026).

The distinction between access and agency also helps explain why AI adoption statistics should be read alongside preparedness indicators. AI Ready ASEAN research reported a notable usage gap in the Philippines: 83% of students had used generative AI for educational purposes, compared with 73% of educators using generative AI for teaching. The presentation uses this gap to raise questions about guidance, policy, and preparedness rather than treating adoption itself as evidence of readiness (ASEAN Foundation, 2026; Tjhi, 2026).

Table 1. Everyday AI Entry Points in Southeast Asia

Domain	Actor	Illustrative activity	Capability focus
Learn	Student	AI Ready ASEAN	Localised AI literacy
Teach	Teacher	CoLearn	Planning and guided tutoring
Trade	Merchant	Gojek	Business recommendations
Care	Health worker	Field scenario	Documentation and decision support
Create	Cultural advocate	BASAibu	Local-language civic participation
Respond	Disaster responder	AHA Centre	Regional coordination

Source: Adapted from Tjhi (2026). The healthcare row is a field scenario, not an attributed programme.

Discussion Points

AI Access Is Not the Same as Agency

The first major discussion point is that access to AI should be understood as a pathway to capability rather than as capability itself. The "equaliser" argument is therefore conditional. Translation can reduce a language barrier; personalised learning can extend educational support; decision tools can reduce the cost of expertise; and generative systems can reduce the cost of creation. Yet each benefit depends on conditions that determine who can actually use the system, understand its limitations, correct its outputs, and carry responsibility when its recommendations are wrong (Tjhi, 2026).

This distinction is particularly important in high-stakes settings. A user who can ask an AI system for information but lacks the skills or authority to challenge an incorrect answer has access without full agency. The same applies where connectivity is intermittent, devices are scarce, or professional support is limited. In these circumstances, the distribution of responsibility matters as much as the distribution of access. An AI deployment that shifts error costs onto the least-resourced users can reproduce the very divide it was intended to narrow.

Ground Truth at the Edge: Four Structural Gaps

A case in point is a remote clinic scenario that can help make a broader point: the hardest AI problem is rarely the model alone (Tjhi, 2026). Four structural gaps appear at the edge - language and voice, cultural fit, infrastructure, and human support - and the same constraints can be visible in both education and healthcare.

The first gap is language and voice. Sparse training data, dialect variation, code-switching, and speech variability can reduce system reliability. The issue is broader than translation: a system that produces fluent language may still fail on local terminology or struggle with voices recorded in noisy, low-bandwidth environments. The second gap is cultural fit. Fluency can mask assumptions that are inappropriate for local institutions, norms, or values, while neighbouring languages can introduce unintended interference. The third gap is infrastructure. Cloud-only systems can be inaccessible in places where bandwidth, electricity, hardware, or device availability is constrained. The fourth gap is human support. Where teachers and clinicians are scarce, the places with the greatest need for decision support may also be the places where oversight is most difficult (Tjhi, 2026).

Table 2. AI Potential, Frictions, and Design Questions

Potential	Typical friction	Design question
On-demand translation	Language and voice gaps	Can users understand and correct the output?
Personalised learning	Skills and supervision	Who guides the learner when the system is wrong?

Potential	Typical friction	Design question
Decision support	Trust, safety, and human support.	Who bears the error and retains final judgment?
Lower-cost creation	Access, data, and cultural fit	Who retains capability after AI assistance?

Source: Author synthesis from Tjhi (2026).

Edge AI as a Design Philosophy

A key argument is on treating edge AI as more than a technical exercise in making a model smaller. It is presented as a design philosophy for local capability: systems should work offline where needed, degrade gracefully when conditions are imperfect, protect data through local processing, fit low-power and lower-cost environments, and broaden the human-in-the-loop by locating capability close to the people who sense the ground first (Tjhi, 2026).

This philosophy turns infrastructure from a background assumption into a core design variable. An AI system intended for Southeast Asian settings must be evaluated not only under ideal cloud-connected conditions, but also under the constraints of the place in which people are expected to use it. The significance of edge AI is therefore institutional as well as technical. When processing, data, and decision support can remain closer to users, communities may have more practical capacity to participate in the AI system rather than simply consume its outputs.

Education, Work, and Trust: Risks with a Local Accent

The second major discussion point is that global AI risks acquire regionally specific forms that may be grouped into education, work, and trust. In education, the concern is cognitive outsourcing: people may increasingly delegate thinking, drafting, or problem-solving to AI without equivalent institutional readiness or pedagogical guidance. In work, the concern is livelihood transition, especially where large service industries depend on tasks highly exposed to AI. In trust, the risk is industrialised deception, in which AI-assisted synthetic media and automated outreach can lower the cost and increase the scale of fraud (Tjhi, 2026).

Table 3. Three Risk Domains Highlighted in the Southeast Asian Context

Domain	Risk framing	Regional evidence / illustration
Education	Cognitive outsourcing	AI Ready ASEAN: 83% of Filipino students vs. 73% of educators reported using generative AI in the cited education context.
Work	Livelihood transition	IMF Working Paper 25/43 identifies Philippine BPO as the sector with the highest proportion of jobs at risk of displacement, while also identifying AI complementarity.
Trust	Industrialised deception	UNODC documents deepfakes, voice cloning, synthetic identities, multilingual chatbots, and automated outreach in Southeast Asian cyber-enabled crime.

Source: ASEAN Foundation (2026), Cucio and Hennig (2025), and UNODC (2025); framing adapted from Tjhi (2026).

The education example illustrates why usage and readiness should not be conflated. The cited AI Ready ASEAN research records higher generative-AI use among Filipino students than among educators. The presentation interprets the difference as a question about guidance, policy, and preparedness. The implication for institutions is not simply to suppress or accelerate usage, but to build the surrounding capability that makes use meaningful: teacher readiness, student literacy, responsible-use norms, and mechanisms for evaluating outputs (ASEAN Foundation, 2026; Tjhi, 2026).

The labour-market example is similarly conditional. Cucio and Hennig (2025) find that the Philippine BPO sector has the highest proportion of jobs at risk of displacement in their sectoral analysis, while also identifying substantial potential for complementarity. This distinction matters

because exposure does not imply a single technological outcome. The policy and management challenge becomes one of transition: identifying which tasks may change, where humans and AI can complement one another, and which skills and institutions are needed for workers to adapt (Cucio & Hennig, 2025).

The trust example shows a different dimension of the same structural problem. UNODC reports that Southeast Asian criminal actors are using AI-enabled tools for deepfakes, synthetic identities, voice cloning, and automated or multilingual interactions that can support fraud and phishing (UNODC, 2025). The scale of the risk demonstrates why trustworthy AI cannot be treated as an after-the-fact technical layer. Identity verification, incident learning, user awareness, organisational controls, and regional cooperation are part of the deployment environment in which AI operates.

From a Model to an Ecosystem: The SEA-LION Journey

One constructive response to these challenges is to build capability together and in the open. SEA-LION is presented as an open story in which regional knowledge, technical capability, and community participation reinforce one another (Tjhi, 2026). AI Singapore's public materials describe SEA-LION as an open multilingual model effort focused on Southeast Asian languages, culture, and use cases, while SEA-HELM provides a regional evaluation layer for linguistic and cultural competencies (AI Singapore, 2025a, 2025b).

The ecosystem view matters because no single model can resolve all of the gaps identified above. In the context of SEA-LION, the ecosystem includes the SEA-LION model family, a developer challenge for applied experimentation, Project SEALD for data and recipes, SEA-HELM for regional evaluation, Project ATLAS as a grassroots data hub and a community platform (Tjhi, 2026).

The ecosystem approach supports a second important conclusion: public-good innovation can create reusable capability. A hard local problem - such as remote care or a low-resource language - can trigger work on data, models, safety, and deployment. Those outputs can then become reusable tools, standards, and people-level skills that spill into services, firms, and institutions. Therefore, public goods are seen not as peripheral charity at the edge of the AI economy, but as capability infrastructure for the wider region (Tjhi, 2026).

An ASEAN AI-for-Good Institution: Sustain, Scale, Steward

One proposal is the establishment of a regional institution for AI-for-good to make the ecosystem logic durable. Its first function is to sustain a technology commons through shared datasets and compute, edge-ready reference systems, open safety and evaluation, and long-term maintenance. Its second function is to scale regional capability through fellowships and exchanges, community labs, deployment partners, and local-language expertise. Its third function is to steward governance through participatory priority-setting, independent evaluation, incident learning, and public-interest procurement (Tjhi, 2026).

Table 4. Proposed Functions of an ASEAN AI-for-Good Institution

Function	Core elements	Purpose in the regional ecosystem
Sustain	Technology commons; shared data and compute; edge-ready systems; open safety; maintenance	Keep regional AI infrastructure available and usable over time.
Scale	Fellowships; exchanges; community labs; deployment partners; local-language expertise	Expand the number of people and organisations able to build and deploy responsibly.
Steward	Participatory priorities; independent evaluation; incident learning; public-interest procurement	Connect technical development with accountable regional governance.

Source: Adapted from Tjhi (2026).

The proposal also has implications for how investment in AI capability is measured. A narrow model-centred view may focus on parameters, benchmark scores, or deployment counts. An ecosystem view adds the less visible infrastructure that allows those models to become useful and accountable:

local data, evaluation practices, trained people, community institutions, deployment partners, and mechanisms for learning from failure. In this sense, capability is not only an object that can be purchased; it is an institutional capacity that can be accumulated and shared.

Collaboration and Democratization as Two Pillars

Two principles sit beneath the proposed positive future: collaboration and democratisation. Collaboration means people of the region build AI and solve problems together. Democratisation means capability and meaningful control are available to those who have the greatest local stakes (Tjhi, 2026). These principles complement the ASEAN governance emphasis on responsible design and the UNESCO emphasis on human dignity and oversight (ASEAN Secretariat, 2024; UNESCO, 2021).

Taken together, the principles suggest that local participation should occur across the AI lifecycle rather than only at the moment of product adoption. Communities can contribute data, articulate needs, test systems, identify failures, develop evaluation standards, and participate in decisions about acceptable uses. This is especially important in low-resource settings because the people closest to the problem often possess context that a centralised technical process cannot fully reconstruct.

Human Flourishing, Subsidiarity, and the Closing Examen

We conclude by citing some parallels with Pope Leo IV's (2026) "Magnifica Humanitas" who offers an image of shared reconstruction:

> *I would like to recall the image of Nehemiah whom we chose as our companion and guide at the outset. Nehemiah heard the cry of a devastated city, brought that pain to prayer, discerned before God, asked for help, received permission to return, organized the work, confronted internal and external resistance and rebuilt the walls of Jerusalem with the assistance of the people, brick by brick. In this era of digital transformation, I see in him a striking parable of our own vocation, which is not to be passive spectators of social and cultural fractures, nor mere commentators on what is crumbling, but men and women prepared to enter the construction sites of history — research laboratories, technology companies, schools, the media, institutions and local communities — in order to rebuild what has collapsed and protect what is threatened. Like Nehemiah, we too are called to unite listening and courage, prayer and responsibility, so that, even when a technocratic mentality or partisan interests seem to prevail, the human city may become a more fitting place to live* (no. 241).

We connect this image to subsidiarity, a principle for locating decisions at the closest competent level to those affected. The concept is useful here not as a technical specification but as a governance principle: regional coordination can coexist with local decision-making when authority is placed near the people who bear the practical consequences.

The closing examination offers five questions that can function as an evaluation frame for future AI initiatives in Southeast Asia: does the system strengthen human dignity; does it expand local capability; can affected communities shape it; does it preserve human judgment; and can the region sustain it together (Tjhi, 2026). These questions translate the article's earlier discussion into a practical checklist that links technical design, institutional capacity, and governance responsibility.

Conclusion

The value of AI in Southeast Asia depends not only on the capabilities of models, but also on the location of knowledge, control, responsibility, and institutional support. AI can lower the cost of expertise and create new opportunities in learning, teaching, commerce, care, civic participation, and disaster response. Yet access remains conditional when connectivity, language, cultural fit, human support, and trust are unevenly distributed.

The discussion of education, work, and trust shows why global AI risks need local interpretation. The cited education usage gap, the Philippine BPO labour-market exposure identified by the IMF, and UNODC reporting on AI-enabled deception each demonstrate a different way in which AI meets existing institutions and vulnerabilities. The response proposed in the keynote is therefore not simply to adopt more AI, but to build capability around AI: edge-ready systems, local-language resources, open evaluation, human oversight, community participation, and durable institutions.

SEA-LION illustrates this ecosystem approach through the combination of open models, regional data work, evaluation, experimentation, and community infrastructure. The proposed ASEAN AI-for-Good institution extends the same logic by assigning responsibilities across sustain, scale, and steward functions. The broader objective is to ensure that the future of AI in Southeast Asia is not merely imported as a finished capability, but built in the region through collaboration and meaningful participation. The article therefore returns to the keynote's final direction: the region's AI future should be built here, together, for the common good (Tjhi, 2026).

References

AI Singapore. (2025a). SEA-LION v4: Our first multimodal release, open, powerful & small. SEA-LION. https://sea-lion.ai/blog/sea-lion-v4-multimodal/

AI Singapore. (2025b). Assessing LLM performance for Southeast Asia: Introducing SEA-HELM. SEA-LION. https://sea-lion.ai/assessing-llm-performance-for-southeast-asia-introducing-sea-helm/

AI Singapore. (2025c). Introducing Aquarium: An open data platform for Southeast Asian languages. SEA-LION. https://sea-lion.ai/blog/aquarium-open-data-platform/

ASEAN Foundation. (2026). AI Ready ASEAN research: Assessing artificial intelligence readiness in ASEAN educational contexts. https://aseanfoundation.org/publication/ai-ready-asean-research-assessing-artificial-intelligence-readiness-in-asean-educational-contexts/

ASEAN Secretariat. (2024). ASEAN guide on AI governance and ethics. ASEAN. https://asean.org/book/asean-guide-on-ai-governance-and-ethics/

ASEAN Secretariat. (2025). Expanded ASEAN guide on AI governance and ethics - Generative AI. ASEAN. https://asean.org/book/expanded-asean-guide-on-ai-governance-and-ethics-generative-ai/

AHA Centre. (n.d.). AHA Centre. https://ahacentre.org/

BASAibu. (n.d.). BASAibu Wiki. https://basaibu.org/

CoLearn. (n.d.). CoLearn Tutor. https://colearn.id/tutor

Cucio, M., & Hennig, T. (2025). Artificial intelligence and the Philippine labor market: Mapping occupational exposure and complementarity (IMF Working Paper No. 25/43). International Monetary Fund. https://doi.org/10.5089/9798229001977.00

Leo XIV. (2026). *Encyclical Letter Magnifica Humanitas of His Holiness Pope Leo XIV on Safeguarding the Human Person in the Time of Artificial Intelligence*. https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html

Tjhi, W. C. (2026). AI for Southeast Asia [Keynote presentation]. International Conference on Economics, Business, and Management Research (ICEBMR) 2026.

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. UNESCO. https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence

United Nations Office on Drugs and Crime. (2025). Malicious use of AI powering illicit activities. https://www.unodc.org/roseap/uploads/documents/Publications/2025/UNODC_Report_Emerging_threats_The_intersection_of_criminal_and_technological_innovation_in_the_use_of_automation_and_AI.pdf