

Indonesian Language Education and Applied Linguistics Reviews
Volume 2, Issue 1, 2026, 37-50
https://doi.org/10.66272/bzqt5g85

MEDIA AKADEMIKA PUBLISHER
ISSN 3110-3480

Available online at:
https://journal.mediaakademika.id/index.php/informs/article/view/82

From Curriculum Alignment to Curriculum Agility: A Model for AI-Responsive Higher Education

Hilma Suryani
UIN Raden Fatah Palembang
**Corresponding Author:* suryanihilma@gmail.com

ARTICLE INFO

Received:
11 January 2026

Revised:
24 February 2026

Accepted:
15 March 2026

Keywords:
curriculum agility; constructive alignment; generative artificial intelligence; curriculum development; AI governance

ABSTRACT

Higher-education curricula must remain internally aligned while responding to rapid changes in generative artificial intelligence, professional practice, and student learning. This article develops an AI-Responsive Curriculum Agility (AIRCA) model through an integrative conceptual review of 32 DOI-verified publications. The synthesis combines constructive alignment, curriculum agility, AI literacy, assessment, faculty capability, governance, and equity. It identifies a recurrent institutional problem: universities often respond to AI through tool guidance or isolated course innovations while programme outcomes, assessment constructs, and quality-assurance processes remain unchanged. The proposed model addresses that problem through six linked stages: environmental scanning, curriculum diagnosis, competency renewal, course and assessment redesign, ethics-equity-stakeholder review, and monitoring through controlled iteration. Assessment redesign functions as the validity anchor of the cycle, while governance distinguishes rapid course-level adjustment from changes to programme standards. The model also separates AI-augmentable performance from capabilities graduates must independently own, including disciplinary judgement, verification, ethical reasoning, and accountable decision-making. Ten design principles and an implementation matrix translate the model into programme and course decisions. The article contributes by extending constructive alignment with a temporal and organisational capacity for evidence-informed change. For Indonesian and other resource-diverse higher-education systems, it proposes contextual rather than uniform agility, with explicit attention to connectivity, language, disability, privacy, staff workload, and unequal access. Curriculum agility is therefore defined not as constant revision, but as the disciplined capacity to make timely, coherent, equitable, and educationally defensible change.

How to Cite: Suryani, H. (2026). From Curriculum Alignment to Curriculum Agility: A Model for AI-Responsive Higher Education. *Indonesian Journal of Pedagogy and Research Development, 2*(1), 37-50. https://doi.org/10.66272/bzqt5g85

Published by:
Media Akademika Publisher

1. INTRODUCTION

Curriculum alignment has long provided a disciplined way to connect intended learning outcomes, teaching and learning activities, and assessment. In constructive alignment, the outcomes communicate the kinds of understanding and performance that students should develop; learning activities give students opportunities to enact those performances; and assessment makes the quality of achievement visible (Biggs, 1996). Alignment remains indispensable because students cannot reasonably be expected to demonstrate capabilities that are absent from instruction or rewarded only indirectly. It also gives programme leaders a common language for mapping progression across courses, reducing duplication, and ensuring that graduate attributes are assessed at an appropriate level.

Internal consistency, however, is not identical to contemporary relevance. A programme can be carefully aligned around outcomes that no longer represent disciplinary knowledge, professional practice, social expectations, or the conditions under which graduates will work. This distinction has become more visible with generative artificial intelligence (GenAI). General-

purpose and discipline-specific systems can now generate text, images, code, summaries, analyses, simulations, and feedback at a speed that changes the practical meaning of many academic tasks. The educational question is therefore not simply whether students are permitted to use a particular application. It is whether existing outcomes, learning sequences, and assessments still provide valid evidence of what graduates know and can responsibly do.

The scale of the challenge is amplified by uneven adoption. Reviews of artificial intelligence in higher education describe rapidly expanding but methodologically and geographically uneven research, with continuing needs for stronger ethical, contextual, and educational grounding (Bond et al., 2024; Crompton & Burke, 2023; Zawacki-Richter et al., 2019). Studies of institutional policies show that universities have moved from emergency responses and prohibitions towards conditional use, disclosure, and pedagogical integration, but guidance remains heterogeneous across institutions and disciplines (Jin et al., 2025; McDonald et al., 2025; Moorhouse et al., 2023). This heterogeneity is not inherently a weakness: disciplinary purposes differ. It becomes a weakness when local decisions are not linked to programme outcomes, assessment standards, or a transparent institutional rationale.

Two simultaneous curriculum risks follow. The first is internal misalignment. A programme may claim that students develop independent analysis, evidence-based judgement, or professional communication while its assessments reward polished products that can be substantially generated by AI without revealing the learner's reasoning. The second is external obsolescence. A programme may continue to assess tasks exactly as they were performed before AI-mediated work became normal in the profession. The first risk weakens the interpretation of achievement; the second weakens graduate relevance. Treating all AI use as misconduct does not resolve either problem, because it can conceal actual practice and prevent students from learning how to verify, disclose, and take responsibility for AI-supported work.

Recent research clarifies why a curriculum-level response is necessary. Educators report that GenAI should prompt greater attention to higher-order thinking, ethical values, learning processes, authentic performance, and relational teaching (Bower et al., 2024). Assessment reviews similarly argue for the development of self-regulation, responsible learning, integrity, assessment literacy, and institutional policy renewal (Xia et al., 2024). Curriculum studies identify institutional support, faculty knowledge, and future expectations as significant conditions for AI-related curriculum development (Abbasi et al., 2025), while emerging curriculum-agility research emphasises adaptable structures, recurring stakeholder engagement, dynamic content, and continuing staff development (Brink et al., 2025). Together, these findings indicate that adding an AI literacy module is insufficient. Reform must connect programme purposes, course design, assessment, faculty capability, governance, and student access.

The need for connection is also theoretical. Technology is not a neutral object inserted after pedagogy has been designed; pedagogical purposes, methods, technologies, values, and contexts shape one another (Fawns, 2022). When GenAI changes the conditions of performance, it can change what counts as participation, authorship, expertise, and evidence. Yet curricular response cannot be determined by technical capability alone. Educational missions, disciplinary standards, and public responsibilities must remain prior. A medical, engineering, education, or language programme may use similar tools but reach different judgements about which activities may be augmented and which decisions must remain demonstrably human-owned.

This article addresses a conceptual and organisational gap. AI curriculum scholarship frequently emphasises what learners should know about AI, while policy scholarship focuses on acceptable use and curriculum-agility work focuses on an institution's capacity to adapt. These strands are complementary, but they do not by themselves explain how evidence should move from environmental change to a governed programme decision. Nor do they consistently specify how institutions can accelerate revision without producing fragmented initiatives, unstable standards, or new inequities. A useful model must therefore preserve constructive alignment while shortening the evidence-to-action cycle and clarifying decision rights.

Accordingly, the article asks three questions: (1) How can constructive alignment be extended to accommodate recurring evidence about AI-mediated change? (2) Which curriculum decisions are required to translate that evidence into coherent programme and course action? and (3) Which

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

governance, ethical, and equity conditions are necessary to make adaptation educationally defensible? The contribution is an AI-Responsive Curriculum Agility (AIRCA) model comprising a six-stage cycle, ten design principles, and an implementation matrix. Although designed for broad international relevance, the model explicitly addresses conditions common in Indonesian higher education, including uneven connectivity, multilingual learning, variation in institutional capacity, and the need to align rapid course innovation with formal programme and quality-assurance structures.

2. METHODS

2.1. Review design and conceptual orientation

The study used an integrative conceptual review. This approach was selected because the research problem crosses bodies of literature that use different units of analysis: constructive alignment focuses on relationships among outcomes, activities, and assessment; AI-in-education research examines applications and learning consequences; policy studies consider acceptable use and institutional governance; and curriculum agility examines organisational responsiveness. The purpose was not to estimate an intervention effect or claim exhaustive coverage of every GenAI publication. It was to identify concepts and mechanisms that could be integrated into a coherent, testable model for curriculum decision-making.

Constructive alignment served as the normative starting point because any responsive curriculum must still connect what programmes intend, what students experience, and what assessment certifies (Biggs, 1996). Curriculum agility supplied the temporal and organisational problem: how institutions detect change and translate evidence into coordinated action (Brink et al., 2025). Research on AI literacy, teaching, assessment, institutional policy, equity, and faculty development supplied the substantive conditions under which that action becomes necessary. This combination allowed responsiveness to be treated as more than rapid content revision.

2.2. Source identification and selection

The analytical corpus comprised 32 DOI-verified publications. One foundational publication on constructive alignment was retained because it defines the central relationship being extended (Biggs, 1996). The remaining sources were published between 2019 and 2026, with the majority from 2023 onwards. Sources were identified through combinations of the terms curriculum alignment, curriculum agility, generative AI, higher education, curriculum reform, assessment, AI literacy, institutional policy, governance, faculty capability, ethics, and equity. Reference chaining was used to locate conceptually central studies and recent reviews. Bibliographic details and DOI resolution were checked against publisher records.

A publication was included when it contributed at least one explicit mechanism or implication for curriculum decision-making. Eligible contributions therefore included empirical studies, systematic or scoping reviews, conceptual frameworks, policy analyses, and a collective manifesto. Studies concerned only with the technical performance of a single system, without a teaching, assessment, policy, or organisational implication, were excluded. The focused corpus was deliberately diverse in method because the proposed model required evidence about curriculum purpose, enacted practice, institutional capacity, and consequences for learners.

2.3. Analytical procedure

Each publication was entered into a concept matrix with nine fields: curriculum purpose, unit of change, evidence source, review rhythm, stakeholder role, AI-related capability, durable human capability, assessment implication, and governance or equity risk. Analysis proceeded in three passes. The first pass coded explicit recommendations or mechanisms, such as AI literacy development, process-based assessment, faculty professional learning, disclosure, stakeholder consultation, or monitoring. The second pass compared mechanisms at programme, course, task, and institutional levels. The third pass linked mechanisms into an ordered cycle by asking what evidence each decision requires and what output it provides to the next decision.

Model construction used four tests. Completeness asked whether the model covered the pathway from an external change to evaluation after implementation. Alignment asked whether each proposed change could be traced to outcomes, learning experiences, and assessment.

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

Governability asked whether programme and course responsibilities were distinguishable. Equity asked whether access, language, disability, privacy, and staff capacity were treated as design conditions rather than later corrections. Stages were revised until each produced a defined decision or artefact for the next stage.

2.4. Scope and interpretive limitations

The review is conceptual rather than systematic in the PRISMA sense. Its value lies in theoretical integration and practical specification, not in a claim that the corpus represents every publication on GenAI and higher education. The field also changes quickly; tool-specific evidence can become outdated even when broader curriculum principles remain relevant. For this reason, the model is tool-agnostic and focuses on changing conditions of performance. It should be empirically examined across disciplines, institutional types, and national systems before claims are made about comparative effectiveness.

Table 1. Analytical dimensions for AI-responsive curriculum review

Dimension	Evidence considered	Curriculum question
Environmental change	AI capability, professional practice, regulation, disciplinary knowledge, student use	What has changed outside and inside the programme?
Curriculum diagnosis	Programme outcomes, course maps, assessment distribution, graduate evidence	Where is the curriculum outdated, duplicated, or vulnerable?
Competency renewal	AI-augmentable tasks, durable human capabilities, AI literacy, ethics	Which capabilities should be added, strengthened, or protected?
Course adaptation	Learning outcomes, content, sequence, pedagogy, and declared tool roles	What can change rapidly without losing programme coherence?
Assessment redesign	Construct, authenticity, process evidence, verification, independent performance	How will assessment remain valid in an AI-rich environment?
Governance and equity	Decision rights, staff capability, access, language, disability, privacy, review data	Who decides, who is supported, and who may be disadvantaged?

3. RESULTS

The synthesis produced a distinction between alignment-dominant and agility-informed curriculum design. The distinction is analytical rather than oppositional. Alignment-dominant design asks whether existing components are coherent; agility-informed design also asks whether those components remain defensible under changing conditions. The AIRCA model therefore places a recurring evidence-and-decision cycle around constructive alignment. It does not make outcomes provisional at every moment. Instead, it creates scheduled opportunities to examine whether the assumptions behind outcomes, pedagogy, and assessment still hold.

Table 1 summarises the analytical dimensions that connect environmental evidence to curriculum questions. The central movement is from description to decision. For example, evidence that students use GenAI for drafting is not itself a curriculum conclusion. It becomes meaningful only when compared with the intended writing capability, the role of drafting in the learning sequence, the assessment construct, and expectations in professional practice. Likewise, evidence that employers use AI does not automatically justify unrestricted use in every academic task. Programmes must decide which uses represent authentic augmentation and which would bypass capabilities graduates remain accountable for.

This reasoning introduces the idea of curriculum half-life: the period during which the assumptions connecting an outcome to professional or disciplinary practice remain sufficiently stable. The half-life is not the same across a programme. Foundational concepts, ethical

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

commitments, and forms of disciplinary reasoning may remain durable, while particular production tasks or technical procedures can change quickly. A curriculum-agility process should therefore avoid subjecting every component to the same review rhythm. Short-cycle review is most useful for tool roles, examples, formative activities, and vulnerable tasks; longer-cycle governance remains appropriate for programme purposes and standards unless strong evidence demands renewal.

The model further distinguishes responsiveness from acceleration. Faster decision-making is valuable only when the decision is connected to valid evidence and authorised responsibility. A hurried requirement that every course use GenAI may be responsive to institutional fashion but unresponsive to learning. Conversely, a carefully bounded pilot completed within one semester may be genuinely agile because it addresses a diagnosed problem, contains safeguards, and generates evidence for programme review. Agility is therefore a quality of the evidence-decision relationship, not the visible frequency of innovation.

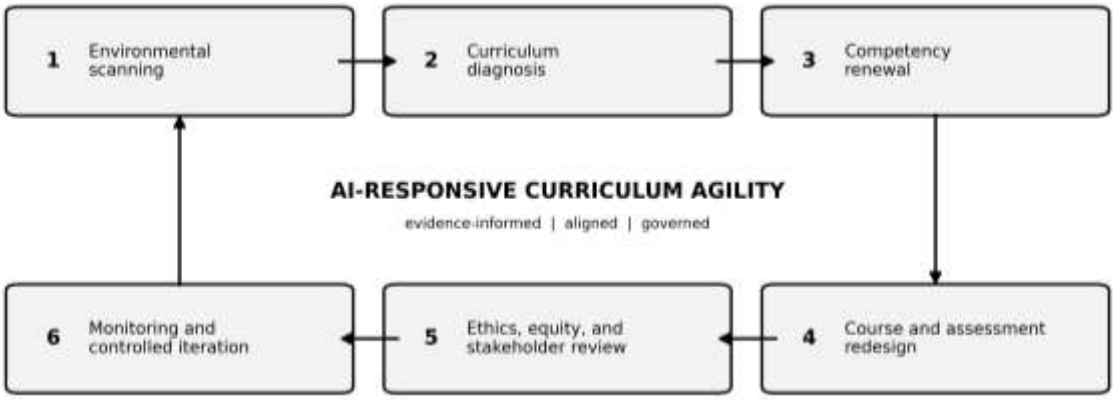


Figure 1. Six-stage cycle for AI-responsive curriculum agility.

3.1. Stage 1: Environmental scanning

Environmental scanning is a disciplined search for changes that may alter programme purposes, knowledge bases, or performance conditions. It includes developments in professional practice, disciplinary research, AI capability, regulation, student behaviour, public expectations, and institutional infrastructure. Scanning is not a technology-news exercise and should not be driven by novelty. The relevant question is whether a change creates a credible implication for what graduates need to know, what they need to do, or how their performance should be interpreted.

Evidence can be assembled through professional standards, tracer studies, employer and alumni consultation, faculty research, student AI-use patterns, assessment samples, integrity cases, and documented changes in workplace tasks. Institutional policy research shows that universities are already producing substantial guidance, sample syllabi, and classroom activities, but that the balance between encouragement and restriction varies (Jin et al., 2025; McDonald et al., 2025). Environmental scanning should therefore compare external signals with local evidence rather than copying the policy of a more highly resourced institution.

A practical output is an environmental change brief. For each development, the brief records the affected discipline or course, the likely duration of the change, the quality of evidence, the potential curriculum consequence, and the urgency of review. High-uncertainty signals can be monitored without triggering immediate reform. High-consequence signals, such as a change in professional accreditation or the routine automation of a core workplace task, should proceed to curriculum diagnosis. This triage prevents both institutional inertia and reactive innovation.

3.2. Stage 2: Curriculum diagnosis

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

Curriculum diagnosis maps the change brief against the enacted programme. Programme outcomes, course outcomes, learning activities, assessment tasks, rubrics, feedback processes, and progression across levels are reviewed for gaps, redundancy, and construct vulnerability. The enacted curriculum matters because formal documents can describe coherence that students do not experience. Diagnosis therefore combines curriculum maps with samples of tasks, student work, feedback, and faculty accounts of how requirements are interpreted.

Construct vulnerability occurs when an assessment can receive a high mark without providing adequate evidence of the capability it is intended to represent. A conventional unsupervised essay may still be appropriate if the construct includes planning, iterative drafting, responsible tool use, source verification, and reflective justification. It becomes vulnerable when only the final prose is observed and substantial reasoning can be delegated without disclosure. Research on institutional assessment guidelines shows why permissions alone are insufficient: acceptable-use statements need to be connected to task design and the evidence required for judgement (Luo, 2024; Moorhouse et al., 2023).

Diagnosis should also identify missed opportunities. Some programmes may overprotect traditional tasks and fail to teach students how AI is actually used in their field. Others may distribute AI activities across many courses without a progression from foundational literacy to discipline-specific application and critical evaluation. An AI curriculum map can mark where students learn about system limitations, practise prompting or workflow design, verify output, examine bias and privacy, disclose use, and demonstrate independent performance. The diagnostic output is a prioritised problem statement, not a list of preferred tools.

3.3. Stage 3: Competency renewal

Competency renewal translates diagnosed problems into decisions about graduate capability. The AIRCA model distinguishes three categories. First, AI-augmentable capabilities are performances in which responsible use of AI is part of contemporary practice, such as generating alternatives, checking code, producing preliminary summaries, or testing communication formats. Second, human-owned capabilities are judgements and responsibilities that graduates must be able to explain and defend, such as selecting evidence, validating claims, interpreting uncertainty, protecting confidential information, and accepting professional accountability. Third, independent demonstration capabilities are the subset that programmes require students to perform without assistance at defined points so that competence can be certified.

The distinction avoids two unhelpful extremes. One treats every task that can be automated as educationally obsolete. Yet automation can increase the importance of foundational knowledge because people must recognise plausible errors and decide when outputs are unsafe. The other treats traditional independent production as the only legitimate evidence of learning. That position can make curricula less authentic and prevent students from developing responsible human-AI collaboration. Curriculum renewal should instead specify the purpose, conditions, and accountability of performance.

AI literacy is necessary but should not be reduced to tool operation. Students need a conceptual understanding of how outputs are generated, where bias and fabrication may arise, what data should not be entered, how authorship and intellectual contribution are represented, and how to evaluate outputs against disciplinary criteria. Pedagogical reviews note potential gains in feedback, creativity, autonomy, and learning support, alongside risks of overreliance and cognitive outsourcing (Farrokhnia et al., 2024; Kasneci et al., 2023; Qian, 2025). Experimental evidence on metacognitive laziness reinforces the need to make monitoring, checking, and justification visible in learning outcomes (Fan et al., 2025).

Renewed outcomes should use observable verbs and specify performance conditions. An outcome such as 'understand ethical AI' is too broad for alignment. A stronger formulation may require students to evaluate an AI-supported disciplinary recommendation, verify its evidence, identify privacy and bias risks, document material AI contributions, and defend the final decision. This formulation can guide activities, rubrics, and independent checkpoints. Competency renewal is complete when the programme has prioritised outcomes and defined where each is introduced, developed, and demonstrated.

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

3.4. Stage 4: Course and assessment redesign

Stage 4 translates programme decisions into learning sequences and assessment. Course adaptation may revise outcome wording, examples, content, teaching activities, task instructions, feedback, and the timing of practice. Rapid adaptation is possible when the programme has already defined guardrails. A course coordinator can update an activity or tool role within an approved outcome, whereas a change to the standard of independent performance requires programme-level approval. This division protects coherence while allowing teachers to respond to new evidence.

Learning design should place AI within a purposeful sequence. Students may first attempt a task using their own knowledge, compare their reasoning with AI-generated alternatives, diagnose strengths and errors, revise the output, and explain which decisions remain theirs. Alternatively, an AI output may be used as a flawed case for critique. These designs treat AI as an object of analysis and a participant in a workflow, not as a source whose fluency is mistaken for authority. An entangled perspective is useful here because the tool, task, values, and context jointly shape what students learn (Fawns, 2022).

Assessment redesign is the validity anchor of the model. If AI changes the conditions under which a task is performed, the programme must examine whether the resulting evidence still supports the intended inference. Scoping and educator studies recommend greater attention to authentic performance, higher-order thinking, learning processes, oral or relational components, self-regulation, and responsible use (Bower et al., 2024; Xia et al., 2024). The purpose is not to make every task resistant to AI. It is to design a defensible evidence system.

A balanced assessment programme can combine four forms of evidence. AI-permitted authentic tasks show how students work with contemporary tools. Process evidence, such as decision logs, annotated prompts, drafts, or version histories, shows how a product developed. Verification evidence requires students to test claims, sources, code, calculations, or recommendations against disciplinary standards. Independent demonstration through an oral defence, live problem, supervised task, or focused explanation confirms that essential capability has not been wholly delegated. No single form is sufficient in all disciplines; validity arises from the pattern.

Rubrics must also change. If a task permits AI but rewards only surface quality, the rubric encourages delegation. Criteria can instead assess problem framing, justification of tool choice, quality of verification, interpretation of uncertainty, ethical handling of data, documentation of material contributions, and the learner's capacity to defend the final work. These criteria shift attention from detecting machine-generated language to evaluating accountable performance. They also make expectations clearer for students who are uncertain about what counts as legitimate assistance.

3.5. Stage 5: Ethics, equity, and stakeholder review

Curriculum innovation can be technically feasible and educationally undesirable. Stage 5 therefore reviews whether proposed changes are ethical, equitable, viable, and aligned with stakeholder responsibilities before full implementation. Students, faculty, programme leaders, quality-assurance staff, information-technology and data-protection units, alumni, professional partners, and where relevant disability-support and language-support specialists contribute different forms of knowledge. Participation is not equivalent to allowing every stakeholder to veto a change. It means that consequences are examined by those who understand and bear them.

Equity review examines access to devices, stable connectivity, paid features, language quality, accessibility, and the time required to learn new systems. A course that implicitly assumes continuous access to a premium English-dominant tool may disadvantage students even when the formal instruction says that AI use is optional. Institutions should provide equivalent pathways, avoid grading the quality of a purchased system, and check whether accessibility tools interact safely with the assessment. Faculty workload is also an equity issue: reforms that require extensive individual verification without staffing or time are unlikely to be implemented consistently.

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

Ethical review addresses privacy, data ownership, intellectual property, bias, transparency, academic integrity, and the consequences of erroneous outputs. Policy frameworks increasingly emphasise responsible and human-centred governance rather than simple permission or prohibition (Bozkurt et al., 2024; Chan, 2023; Crompton et al., 2026). The curriculum should translate these values into student action: what information may be shared, when disclosure is required, how sources are checked, how AI contributions are represented, and who remains accountable for a final decision.

The output is a readiness decision with conditions. A proposal may be approved, revised, piloted, or deferred. Conditions can include an alternative task, institutional access, a privacy-approved tool category, faculty development, student orientation, revised instructions, or a limited pilot. Recording the rationale matters because it allows later monitoring to examine whether anticipated benefits and risks occurred. It also prevents ethical and equity concerns from becoming informal, inconsistent exceptions.

3.6. Stage 6: Monitoring and controlled iteration

Agility requires evidence after implementation. Monitoring asks whether the redesigned curriculum improved learning and validity without creating unacceptable burdens or exclusions. Useful indicators include the quality of student reasoning, performance at independent checkpoints, error and verification patterns, disclosure quality, student workload, staff workload, tool access, integrity cases, privacy incidents, and differences across student groups. Measures should be proportionate; the objective is not to produce a surveillance system but to support defensible decisions.

Controlled iteration distinguishes the speed of change from the level of authority. Course-level examples, activities, formative tasks, and instructions can often be adjusted within a semester when evidence shows a clear problem. Changes to programme outcomes, progression standards, high-stakes assessment patterns, or accreditation commitments require scheduled governance. This dual rhythm is central to AIRCA. Without course-level flexibility, evidence arrives faster than action; without programme governance, local innovations accumulate without coherence.

Monitoring completes and restarts the cycle. Unexpected dependence, exclusion, staff burden, or weak learning returns to diagnosis. Evidence of a new professional workflow may return to competency renewal. A successful pilot may be scaled through programme mapping and staff development. The model therefore treats implementation as a source of knowledge rather than the endpoint of curriculum work. Agility is institutional learning with boundaries, not continuous disruption.

3.7. Ten principles of AI-responsive curriculum agility

1. Protect the educational mission before selecting technologies.
2. Distinguish AI-augmentable performance from capabilities graduates must independently own.
3. Preserve constructive alignment while shortening the evidence-to-action cycle.
4. Use programme maps to prevent isolated, duplicative, or contradictory AI initiatives.
5. Redesign assessment whenever AI changes the conditions under which performance is produced.
6. Treat ethics, privacy, accessibility, language, and equity as curriculum design criteria.
7. Involve students, faculty, alumni, professional partners, and governance units in diagnosis and review.
8. Develop faculty curriculum judgement, AI literacy, and assessment literacy before mandating integration.
9. Monitor unintended effects, including cognitive dependence, workload, privacy risk, and unequal access.
10. Use controlled iteration: rapid course adaptation within transparent programme governance.

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

The ten principles summarise constraints that apply across all six stages. Their sequence moves from purpose to capability, coherence, assessment, equity, participation, implementation capacity, and review. They are not a checklist for tool adoption. A programme may follow them and conclude that a particular use should not be implemented. The important condition is that the decision is based on evidence, curriculum purpose, and transparent criteria rather than enthusiasm, fear, or institutional imitation.

Table 2 shows how the agility-informed approach extends rather than rejects alignment. The main addition is a continuing relationship between the curriculum and its environment. This relationship requires institutional capacity: data that can be interpreted, staff who can make curriculum judgements, and governance that can authorise proportionate change. Technology procurement alone cannot create that capacity.

Table 2. From curriculum alignment to curriculum agility

Curriculum feature	Alignment-dominant approach	Agility-informed approach
Primary purpose	Internal consistency among outcomes, teaching, and assessment	Internal consistency plus continuing external relevance
Review rhythm	Periodic programme review	Continuous scanning with scheduled decision points
Change unit	Programme documentation	Programme, course, task, and assessment evidence
Stakeholder input	Consultation during formal review	Recurring evidence, feedback, and co-creation
Outcome stability	Outcomes remain relatively fixed	Outcomes can be renewed through controlled governance
Assessment	Confirms stated outcomes	Also tests construct validity under AI-assisted performance
Staff capability	Implementation of approved curriculum	Ongoing curriculum judgement and AI-aware assessment literacy

3.8. Implementation and accountability matrix

The implementation matrix in Table 3 turns the cycle into a sequence of accountable decisions. Institutions can adapt the evidence sources and committee names without altering the logic. A small programme may complete several responsibilities through one team, while a large university may distribute them across curriculum, academic-integrity, information-security, and quality-assurance units. What should not disappear is the handoff between stages. Tool guidance that does not reach assessment design, or course innovation that does not generate review evidence, leaves the cycle incomplete.

Table 3. Decisions, evidence, and accountability across the AIRCA cycle

Stage	Key decision	Required evidence/output	Primary accountability
1. Scan	Does the change affect programme purposes or performance conditions?	Environmental change brief with urgency and evidence rating	Programme team with external and institutional input
2. Diagnose	Where are gaps, duplication, or construct vulnerabilities?	Curriculum and assessment map; prioritised problem statement	Programme leader and course coordinators
3. Renew	What is augmentable, human-owned, and independently demonstrable?	Revised outcomes and capability-progression map	Programme committee and disciplinary experts
4. Redesign	Which learning sequence and evidence system enact the outcomes?	Revised activities, assessments, rubrics, and disclosure requirements	Course team within programme guardrails
5. Review	Is the change ethical, equitable, feasible, and supported?	Readiness decision, safeguards, alternatives, and development plan	Cross-functional stakeholder review
6. Monitor	What should be retained, revised, scaled, or stopped?	Outcome, validity, workload, access, and risk indicators	Course and programme governance at defined intervals

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

4. DISCUSSIONS

4.1. Extending constructive alignment

The first theoretical contribution is to extend constructive alignment with a temporal and organisational dimension. Alignment traditionally clarifies relationships within a curriculum at a given point. AIRCA asks how those relationships are re-examined when the environment changes. This extension does not weaken outcome-based design; it protects it from becoming internally coherent but externally outdated. Curriculum agility is therefore best understood as the institutional capability to review the assumptions that support alignment and to revise them through evidence-informed governance.

The model also clarifies that responsiveness occurs at several levels. Programme outcomes establish shared purposes and standards. Courses enact progression. Tasks define local performance conditions. Assessment evidence supports decisions about achievement. Institutional policies provide ethical and operational boundaries. AI-related reform becomes fragmented when one level changes without the others. For example, a central policy may permit GenAI while course rubrics continue to reward unexamined products; conversely, an innovative task may require disclosure that is unsupported by institutional privacy or data guidance.

4.2. Assessment validity and accountable performance

The second contribution is to position assessment redesign as the validity anchor of AI-responsive curriculum. Public discussion often frames assessment as a problem of detecting generated text. That framing is too narrow because detection does not establish what students have learned, and unreliable detection can create unfair accusations. The curriculum question is whether the assessment evidence supports the intended inference. When performance is AI-assisted, validity can be strengthened by observing decisions, verification, explanation, and independent capability rather than relying on the apparent authorship of a single product.

This approach is compatible with authenticity. Authenticity does not mean allowing every available tool in every task; it means reproducing the purposes, constraints, responsibilities, and quality standards of a meaningful practice. In professional work, AI use may be expected, but confidential data, safety-critical decisions, and final accountability are constrained. An authentic assessment can therefore permit assistance for ideation or drafting while requiring independent verification and oral defence. The resulting evidence is closer to responsible practice than either unrestricted delegation or an artificial assumption that the technology does not exist.

A programme-level evidence system is especially important. If every course responds by adding an oral defence, student and staff workload may become unsustainable. If no course includes independent demonstration, the programme may be unable to certify foundational capability. Curriculum mapping allows different forms of evidence to be distributed strategically across a programme. It also allows later stages to increase complexity: students can move from recognising limitations, to verifying outputs, to designing responsible workflows, and finally to defending consequential decisions.

4.3. Governance without curricular paralysis

The third contribution is the concept of controlled iteration. Universities often experience a mismatch between the speed of tool change and the pace of formal curriculum approval. A demand that every small adaptation pass through full programme review produces paralysis, while unrestricted local experimentation produces inconsistency. AIRCA resolves the tension by linking change authority to consequence. Course teams can modify examples, scaffolds, formative activities, and tool categories within approved outcomes and safeguards. Programme governance retains authority over graduate outcomes, high-stakes assessment architecture, progression, and standards.

Global Delphi evidence indicates that effective GenAI governance requires a balance of policy, capacity building, shared values, and ongoing review (Crompton et al., 2026). The model translates that balance into curriculum work. Governance is not a final compliance checkpoint; it specifies

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

who interprets evidence, who authorises a change, what safeguards are required, and when the decision will be revisited. Transparent decision records also support consistency across programmes without assuming that all disciplines should adopt identical practices.

4.4. Faculty capability as curriculum infrastructure

Faculty capability is not an implementation detail. It is curriculum infrastructure. Staff must understand enough about AI to judge its relevance and limitations, but technical fluency alone is insufficient. They also need assessment literacy, curriculum-mapping skills, ethical and privacy awareness, and the confidence to discuss uncertain practices with students. Faculty studies report varied profiles, self-efficacy, and professional-development needs, suggesting that uniform training will not meet all starting points (Mah & Groß, 2024).

Professional learning should therefore be organised around curriculum decisions rather than product demonstrations. Staff can audit an existing task, identify its construct, test how AI changes the evidence, design a verification or independent component, and examine equity implications. This produces an immediately usable revision and develops judgement that can transfer when tools change. Communities of practice can compare disciplinary solutions and reduce duplicated effort, while programme maps show where local changes contribute to shared outcomes.

4.5. Contextualising agility in Indonesian higher education

For Indonesian higher education, a contextual interpretation is crucial. Institutions differ in location, connectivity, funding, language ecology, discipline mix, staff workload, and access to licensed systems. An AI-responsive curriculum cannot assume that all students have stable broadband, premium subscriptions, current devices, or equivalent English proficiency. Nor can it assume that a policy written for a research-intensive metropolitan university can be transferred unchanged to a smaller or regionally located institution. The common element should be the decision process, not a mandatory tool.

AIRCA supports contextualisation by separating educational requirements from technological means. A programme may require students to compare, verify, document, and defend AI-supported work, while providing different approved tools or a non-AI alternative. Environmental scanning can include local employer practice and Indonesian regulatory guidance alongside international evidence. Stakeholder review can examine Bahasa Indonesia and local-language performance, culturally situated bias, and the effect of English-dominant outputs on disciplinary voice. Monitoring can disaggregate access and achievement rather than interpreting average performance as evidence of fairness.

The model can also align with existing quality-assurance routines. Environmental briefs and curriculum maps can inform programme evaluation; renewed outcomes can be connected to graduate profiles; redesigned assessments can be documented through course plans and rubrics; and monitoring evidence can feed annual review. Embedding agility in familiar processes reduces the risk that AI reform becomes a temporary project owned only by enthusiastic individuals.

4.6. Research implications and limitations

The AIRCA model generates an empirical research agenda. Comparative case studies can examine how institutions move evidence across the six stages and where decisions stall. Design-based studies can test whether combined AI-permitted, process, verification, and independent evidence improves the interpretability of assessment. Longitudinal research can investigate whether controlled iteration increases decision speed without increasing curricular fragmentation. Equity research should examine who benefits from AI-integrated curricula, who bears additional costs, and whether alternative pathways produce comparable opportunities to demonstrate learning.

Several limitations remain. The conceptual review was focused rather than exhaustive, and the emerging evidence base contains a large proportion of perception studies and rapidly changing tool contexts. The model has not yet been validated as a complete institutional cycle. Its six stages are presented sequentially for clarity, but actual curriculum work may move backward when new evidence or constraints emerge. Finally, the distinction between augmentable and

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

human-owned capability is necessarily normative and disciplinary; the model structures that judgement but cannot make it on behalf of a programme.

5. CONCLUSION

AI-responsive curriculum development requires more than adding content, issuing tool rules, or redesigning isolated assignments. It requires an institutional capacity to detect consequential change, diagnose curriculum effects, renew capabilities, redesign learning and assessment, review ethics and equity, and learn from implementation. The AIRCA model integrates these responsibilities in a six-stage cycle while retaining constructive alignment as the basis of educational coherence.

The model's central claim is that agility and stability are not opposites. Stable educational missions and programme standards make rapid local adaptation meaningful; recurring evidence protects those standards from becoming obsolete. Assessment redesign anchors validity, competency renewal distinguishes augmentation from accountable human capability, and controlled iteration connects course flexibility with programme governance. Ethics, privacy, accessibility, language, and equity operate as design criteria throughout rather than as final checks.

Curriculum agility should therefore not be measured by how frequently a university changes its courses or how many AI tools it adopts. It is the disciplined ability to make timely, coherent, evidence-informed, and educationally defensible decisions. That capability offers a more durable response to GenAI than any tool-specific policy because it enables institutions to continue learning as technology, professions, and student practices change.

6. REFERENCES

Abbasi, B. N., Wu, Y., & Luo, Z. (2025). Exploring the impact of artificial intelligence on curriculum development in global higher education institutions. *Education and Information Technologies*, *30*, 547–581. https://doi.org/10.1007/s10639-024-13113-z

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, *32*, 347–364. https://doi.org/10.1007/BF00138871

Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education*, *21*, 4. https://doi.org/10.1186/s41239-023-00436-z

Bower, M., Torrington, J., Lai, J. W. M., Petocz, P., & Alfano, M. (2024). How should we change teaching and assessment in response to increasingly powerful generative artificial intelligence? Outcomes of the ChatGPT teacher survey. *Education and Information Technologies*, *29*, 15403–15439. https://doi.org/10.1007/s10639-023-12405-0

Bozkurt, A., Xiao, J., Farrow, R., Bai, J. Y. H., Nerantzi, C., Moore, S., Dron, J., et al. (2024). The manifesto for teaching and learning in a time of generative AI: A critical collective stance to better navigate the future. *Open Praxis*, *16*(4), 487–513. https://doi.org/10.55982/openpraxis.16.4.777

Brink, S. C., de Hei, M., Sjoer, E., Carlsson, C. J., Georgsson, F., Keller, E., McCartan, C., Enelund, M., Lyng, R., & Admiraal, W. (2025). Curriculum agility principles for transformative innovation in engineering education. *European Journal of Engineering Education*, *50*(3), 455–471. https://doi.org/10.1080/03043797.2024.2398165

Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, *20*, 38. https://doi.org/10.1186/s41239-023-00408-3

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

Chu, T. S., & Ashraf, M. (2025). Artificial intelligence in curriculum design: A data-driven approach to higher education innovation. *Knowledge*, *5*(3), 14. https://doi.org/10.3390/knowledge5030014

Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating? Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*(2), 228–239. https://doi.org/10.1080/14703297.2023.2190148

Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, *20*, 22. https://doi.org/10.1186/s41239-023-00392-8

Crompton, H., Burke, D., Nickel, C., Bozkurt, A., Miao, F., Sharples, M., et al. (2026). Governing generative AI in higher education: A global Delphi study on policy and practice. *International Journal of Educational Technology in Higher Education*, *23*, 21. https://doi.org/10.1186/s41239-026-00602-z

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, *71*, 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

Eaton, S. E. (2023). Postplagiarism: Transdisciplinary ethics and integrity in the age of artificial intelligence and neurotechnology. *International Journal for Educational Integrity*, *19*, 23. https://doi.org/10.1007/s40979-023-00144-1

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, *56*(2), 489–530. https://doi.org/10.1111/bjet.13544

Farrokhnia, M., Banihashem, S. K., Noroozi, O., & Wals, A. (2024). A SWOT analysis of ChatGPT: Implications for educational practice and research. *Innovations in Education and Teaching International*, *61*(3), 460–474. https://doi.org/10.1080/14703297.2023.2195846

Fawns, T. (2022). An entangled pedagogy: Looking beyond the pedagogy-technology dichotomy. *Postdigital Science and Education*, *4*, 711–728. https://doi.org/10.1007/s42438-022-00302-7

Ifenthaler, D., Majumdar, R., Gorissen, P., Judge, M., Mishra, S., Raffaghelli, J., & Shimada, A. (2024). Artificial intelligence in education: Implications for policymakers, researchers, and practitioners. *Technology, Knowledge and Learning*, *29*, 1693–1710. https://doi.org/10.1007/s10758-024-09747-0

Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence*, *8*, 100348. https://doi.org/10.1016/j.caeai.2024.100348

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., et al. (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

Lee, D., Arnold, M., Srivastava, A., Plastow, K., Strelan, P., Ploeckl, F., Lekkas, D., & Palmer, E. (2024). The impact of generative AI on higher education learning and teaching: A study of educators' perspectives. *Computers and Education: Artificial Intelligence*, *6*, 100221. https://doi.org/10.1016/j.caeai.2024.100221

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.

Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences*, *13*(4), 410. https://doi.org/10.3390/educsci13040410

Luo, J. (2024). A critical review of GenAI policies in higher education assessment: A call to reconsider the originality of students' work. *Assessment & Evaluation in Higher Education*, *49*(5), 651–664. https://doi.org/10.1080/02602938.2024.2309963

Ma, Y., Su, Y., Li, M., Zhang, Y., Chai, W., Huang, A., & Zhao, X. (2025). Preparing students for an AI-driven world: Generative AI and curriculum reform in higher education. *Frontiers of Digital Education*, *2*, 30. https://doi.org/10.1007/s44366-025-0067-6

Mah, D. K., & Groß, N. (2024). Artificial intelligence in higher education: Exploring faculty use, self-efficacy, distinct profiles, and professional development needs. *International Journal of Educational Technology in Higher Education*, *21*, 58. https://doi.org/10.1186/s41239-024-00490-1

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

Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open*, *5*, 100151. https://doi.org/10.1016/j.caeo.2023.100151

Qian, Y. (2025). Pedagogical applications of generative AI in higher education: A systematic review of the field. *TechTrends*, *69*, 1105–1120. https://doi.org/10.1007/s11528-025-01100-1

Southworth, J., Migliaccio, K., Glover, J., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for AI across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, *4*, 100127. https://doi.org/10.1016/j.caeai.2023.100127

Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, *10*, 15. https://doi.org/10.1186/s40561-023-00237-x

Walter, Y. (2024). Embracing the future of artificial intelligence in the classroom: The relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education*, *21*, 15. https://doi.org/10.1186/s41239-024-00448-3

Xia, Q., Weng, X., Ouyang, F., Lin, T. J., & Chiu, T. K. F. (2024). A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education*, *21*, 40. https://doi.org/10.1186/s41239-024-00468-z

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education*, *16*, 39. https://doi.org/10.1186/s41239-019-0171-0

Copyright (c) 2026 Suryani
This work is licensed under a Creative Commons Attribution 4.0 International License.