EDUCATION LEADERSHIP WHITEPAPER

Beyond Efficiency

Artificial Intelligence and the Formation of Persons in University Education

A White Paper for Singapore and Asia from an Evangelical Christian perspective

"AI should extend human capability without displacing the intellectual effort, judgment, relationships, and responsibility that education exists to cultivate"

Hoong Chuin Lau

Updated 8 August 2026

For professors, academic leaders, and senior administrators in universities

The authors acknowledge the use of OpenAI's ChatGPT during the preparation of this manuscript. The system was used to assist with brainstorming, improving language clarity, and drafting portions of the text.

Abstract

Generative artificial intelligence has entered university education primarily through questions of productivity, assessment, academic integrity, and institutional competitiveness. These questions are important but incomplete. AI is not merely a tool used within education; it is a formative technology that can reshape the practices through which students learn to attend, reason, write, judge, relate, and accept responsibility. Its educational value therefore cannot be inferred from improved task performance alone. An AI system may help a student produce a stronger answer while either strengthening or weakening the student's capacity to understand, evaluate, and produce such an answer independently.

This paper develops a Protestant (covering both evangelical and Pentecostal) Christian framework for evaluating and governing AI in university education. It places emerging empirical evidence in dialogue with theological accounts of the *imago Dei*, creaturely limitation, the fall and idolatry, the cultural mandate and common grace, embodied and relational knowing, the person and work of Christ, the Spirit's illumination and discernment, truth and wisdom, vocation, stewardship, community, grace, and hope. Developed with particular attention to Asian and Chinese educational contexts, the framework places Christian theology in critical dialogue with Confucian accounts of learning, self-cultivation, relational responsibility, and harmony, while refusing to treat Asia or Chinese culture as homogeneous. It proposes **formation-sensitive augmentation** as a governing principle: *AI should augment learning without displacing the intellectual effort, authorship, judgment, relationships, and moral responsibility that university education exists to cultivate*. The framework distinguishes six domains of impact, namely, epistemic, cognitive, moral, relational, vocational, and institutional, and translates them into pedagogical and governance questions for professors and university leaders. The paper rejects both technological inevitability and indiscriminate prohibition. It argues instead that faithful institutional leadership requires universities to specify which human capacities their curricula are intended to form, identify where AI use supports or bypasses those capacities, and redesign teaching, assessment, faculty work, and institutional policy accordingly.

Keywords: artificial intelligence; generative AI; higher education; Christian education; Christian theology; Confucian education; Asia; Chinese universities; human formation; university leadership

Executive summary

The central challenge posed by generative AI is not that machines can now complete many academic tasks. It is that universities may mistake the successful completion of those tasks for the education of the person completing them.

The paper therefore advances five propositions:

1. **Educational performance and human learning are not equivalent.** AI can raise the quality or speed of an output while reducing the cognitive work through which durable understanding develops.
2. **University education is formative before it is transactional.** It forms habits of attention, truthfulness, judgment, responsibility, cooperation, and service, whether institutions acknowledge this formation or not.
3. **AI use is morally and pedagogically differentiated.** Its effects depend on the learner's stage of development, the capability being formed, the task being delegated, the design of the system, and the accountability structure surrounding its use.
4. **Human responsibility cannot be delegated with task execution.** Students, professors, and institutions remain accountable for claims, decisions, designs, and consequences produced with AI assistance.
5. **Universities need a positive account of educational purpose before they can govern AI well.** Rules about permitted tools will remain unstable unless they are grounded in explicit judgments about the human capacities, relationships, and goods that education should cultivate.

The paper calls this approach **formation-sensitive augmentation**. The framework does not ask whether AI is simply allowed. It asks what an activity is intended to form; which part of the work may be delegated; what understanding, judgment, relationship, and responsibility must remain humanly retained; and how the university will know whether genuine learning has occurred. It recommends four context-dependent postures—require, permit, restrict or stage, and prohibit for a specified activity—rather than one institution-wide answer.

For professors, the immediate priority is to distinguish assistance-rich learning from credible assurance of competence. Courses should progress, where appropriate, from protected foundational practice to constrained support and then to governed use of professional tools. Students should learn to verify, revise, disclose, explain, and remain answerable for AI-supported work. For university leaders, the priority is to align institutional strategy, programme outcomes, assessment, procurement, faculty development, data governance, and evaluation around educational purpose rather than adoption targets.

In Asian (and especially Chinese) contexts, this judgment must account for both the resources and distortions of existing educational cultures. Disciplined learning, respect for teachers, family obligation, relational responsibility, and harmony can resist shallow individualism. They can also become overwork, deference, performance identity, shame, and conformity. AI can widen multilingual access and reduce the fear of asking questions, yet it can also intensify credential competition, privatize learning, and amplify confident authority. Institutions should therefore conduct local, participatory evaluation rather than infer effects from a generic image of "Asian learners."

Content

1. Introduction: when performance can be outsourced

The arrival of generative artificial intelligence in university education has produced an understandable rush toward policy. Universities have revised academic-integrity codes, redesigned assessments, licensed AI platforms, created faculty guidelines, and debated whether students should be encouraged, restricted, or required to use these systems. Much of this activity has been necessary. Yet policy has often moved faster than educational reflection. Institutions are deciding where AI may be used before reaching sufficient clarity about what education is for and what kinds of human development its practices are meant to produce.

This is not simply a temporary governance problem caused by a new technology. Generative AI exposes a longstanding ambiguity within the modern university. Is university education principally the delivery of information, the certification of competent performance, preparation for employment, the formation of intellectual and moral agency, or some combination of these goods? When students can generate explanations, code, analyses, designs, and polished prose within seconds, the distinction becomes unavoidable. If the purpose of an assignment is merely to obtain a competent artifact, AI may complete it efficiently. If the purpose is to form the student who struggles to understand, judge, create, revise, and take responsibility for that artifact, efficiency may be an inadequate—and sometimes misleading—measure of success.

The empirical literature already warns against simple conclusions. Carefully designed AI tutors can provide adaptive explanations, rapid feedback, and learning support at a scale difficult for individual instructors to match. In some controlled settings they have improved immediate learning outcomes and student engagement. At the same time, unguarded access to powerful language models can improve performance during assisted practice while impairing subsequent independent performance. Research on writing, knowledge work, and conversational sycophancy raises further concerns about cognitive offloading, shallow ownership, uncritical trust, and the reinforcement of users' preferred judgments. The evidence is still young, heterogeneous, and frequently limited to short-term outcomes. It does not justify a verdict that AI either improves education or inevitably weakens it. It supports a more demanding conclusion: educational effects depend on what is delegated, by whom, at what stage of learning, through what system design, and toward what conception of educational purpose.

For Christian professors and university leaders, this is an anthropological and theological question as well as an empirical one. Technologies are not neutral containers into which human beings pour their purposes. Their affordances make some actions easier, normalize some relationships, distribute power in particular ways, and encourage certain habits over others. Repeated practices shape attention and desire. A system designed to answer immediately may cultivate intellectual dependence or, under different conditions, provide the feedback that allows a learner to persist. A conversational agent that is endlessly available and affirming may support an isolated student, yet also displace the vulnerability and correction found in human relationships. A writing assistant may help a novice see the structure of a

good argument, or it may allow that novice to bypass the difficult process through which writing becomes a mode of thought.

A Christian account begins with neither fear of technology nor confidence in technological progress. Human beings are creatures made in the image of God, endowed with dignity and called to cultivate creation responsibly. Human making can be an expression of creativity, stewardship, and love of neighbor. Yet human beings are also finite and fallen. Technologies can amplify service and repair, but also pride, domination, deception, passivity, and the desire to escape creaturely limits. The Christian tradition therefore resists two rival errors: treating AI as intrinsically corrupting, and treating technological adoption as an inevitable path to human flourishing.

Theologically, the university student is more than a producer of outputs or a future unit of labor. The student is a person whose worth is not earned through performance, whose intellectual gifts carry responsibilities, and whose education should be ordered toward truthful understanding and the service of others. Learning is not merely information acquisition. It involves patience, intellectual humility, disciplined attention, embodied practice, correction by teachers and peers, and the gradual acquisition of judgment. These goods are not all made more valuable by being made faster. Some arise precisely through sustained effort, dependence, failure, and revision.

This paper proposes **formation-sensitive augmentation** as a framework for the use and governance of AI in university education. The term “augmentation” recognizes that AI can genuinely extend human capability, widen access to explanation, reduce avoidable administrative labor, and enable new forms of inquiry. “Formation-sensitive” insists that such benefits be evaluated against the capacities and habits a particular educational practice is meant to cultivate. The governing question is not whether AI was used, but whether its use supported or bypassed the learner’s development into a truthful, capable, responsible, and relational agent.

This approach moves beyond the binary distinction between permission and prohibition. A use that is appropriate for an expert may be damaging for a novice who has not yet acquired the underlying mental model. Assistance that supports formative feedback may be inappropriate in a summative assessment intended to establish independent competence. Delegating transcription or routine formatting differs from delegating the interpretation of evidence, the moral evaluation of a design, or responsibility for a consequential decision. The same tool can therefore be educationally constructive in one context and deformative in another.

The argument is addressed to two overlapping constituencies. Leaders of Christian universities bear responsibility for aligning institutional AI strategy with their universities’ theological missions rather than merely adding religious language to policies inherited from the technology sector. Christian professors and leaders serving in secular universities cannot presume shared doctrinal commitments, but they can articulate widely intelligible educational goods: truthfulness, intellectual agency, responsible authorship, disciplined inquiry, human dignity, relational learning, professional accountability, and the common good. The Christian contribution in a plural institution should not be covertly sectarian. It should be a

transparent, reasoned account of why these goods matter and how AI-mediated practices may strengthen or erode them.

The paper proceeds in four movements. First, it clarifies the purposes of university education and reviews what current evidence does and does not establish about generative AI's educational effects. Second, it develops a Protestant Christian anthropology of learning and technology. Third, it examines AI's formative effects across epistemic, cognitive, moral, relational, vocational, and institutional domains. Finally, it translates the analysis into a decision framework and recommendations for course design, assessment, faculty practice, institutional governance, and future research.

The burden of the argument is constructive. Universities should neither preserve every inherited practice nor automate every practice that can be automated. They should ask which forms of human attention, effort, judgment, presence, and responsibility remain essential to their educational mission—and then design the use of AI around the formation of those goods.

2. What university education is for

Generative AI does not create the university's uncertainty about its purpose; it exposes it. When an AI system can produce an acceptable essay, explanation, design, or program in seconds, professors must decide whether the artifact was the point of the exercise or evidence of a human capability that the exercise was intended to develop. University leaders face the same question at institutional scale. If education is treated primarily as the efficient production of credentials, employable skills, research outputs, and satisfied customers, extensive automation appears self-evidently desirable. If education also concerns the formation of people capable of seeking truth, exercising judgment, acting responsibly, serving others, and participating in communities of inquiry, the evaluation becomes more demanding.

No single formula can encompass every university. Research-intensive and teaching-intensive institutions differ; so do liberal-arts colleges, technological universities, professional schools, and explicitly Christian institutions. Their legitimate purposes are plural. The argument here is therefore not that employability, technical competence, economic contribution, or personal advancement are unworthy aims. Each is a genuine good. The problem arises when these goods become detached from the truth of what students know, the quality of their judgment, the integrity of their conduct, the ends toward which their capabilities are directed, and their accountability to other people.

This paper proposes an **ordered plurality of educational goods**. University education should pursue at least five mutually supporting purposes:

6. the pursuit, testing, integration, and responsible communication of truth;
7. the development of intellectual capability and practical judgment;
8. the cultivation of dispositions appropriate to truthful and responsible inquiry;
9. preparation for vocation, work, and service to the common good; and
10. participation in communities of knowledge, disagreement, apprenticeship, and accountability.

These purposes cannot be reduced to one another. Employment without judgment can produce technically capable but irresponsible professionals. Character without disciplinary competence can become well-intentioned ineffectiveness. Knowledge without love of neighbor can become domination. Community without truth can become conformity. The ordering matters: productivity and performance are educationally valuable when they serve truthful knowledge, responsible agency, and human flourishing, but they are inadequate as final measures of education.

2.1 Truth as a constitutive good

Universities do many things, but the disciplined pursuit of truth is constitutive of their public identity. Teaching invites students into bodies of knowledge that precede them. Research subjects inherited claims to scrutiny and extends what can responsibly be known. Assessment certifies that particular claims about a student's understanding or competence are warranted. Even disagreement about truth presupposes norms of evidence, intelligibility, criticism, and correction. A university that becomes indifferent to whether its claims are true does not merely perform its function poorly; it compromises the basis on which society trusts it.

A Christian account deepens this commitment without implying that Christian scholars possess truth infallibly. Scripture presents wisdom as beginning in a right relation to God (Prov. 1:7), commands truthful speech (Eph. 4:25), and joins knowledge to love rather than pride (1 Cor. 8:1). Because creation is intelligible and belongs to God, inquiry can be received as a vocation rather than a struggle for possession. Because human beings are finite and fallen, Christian scholarship should also be marked by epistemic humility, openness to correction, and suspicion of self-serving certainty. Arthur Holmes's account of Christian higher education resists separating faith from learning, while Nicholas Wolterstorff places scholarship within the wider horizon of *shalom*, justice, and responsible action rather than treating knowledge as an isolated possession (Holmes 1987; Wolterstorff 2004).

Within a plural university, the theological warrant need not be concealed, but institutional arguments should not depend on colleagues first accepting Christian doctrine. The shared claim is more modest and publicly testable: students should learn to distinguish assertion from evidence, fluency from understanding, consensus from truth, and confidence from warranted judgment. They should be able to identify the provenance and limitations of a claim, expose it to criticism, revise it when necessary, and communicate it without deception. These are not sectarian impositions. They are conditions of trustworthy scholarship and professional life.

AI makes these practices more important because language models can generate coherent statements without possessing the human relation to truth for which students and scholars remain accountable. The central issue is not merely hallucination. Even an accurate AI-generated answer can bypass the learner's encounter with evidence and reasons. Conversely, an AI-produced claim can become educationally valuable when students must trace its sources, reconstruct its reasoning, compare it with alternatives, identify uncertainty, and decide whether it should be affirmed. The relevant educational outcome is therefore not simply a true-looking output, but a person increasingly capable of seeking, recognizing, testing, and speaking truth.

2.2 Intellectual agency and practical judgment

University education should develop more than the ability to reproduce information. It should enlarge **intellectual agency**: the capacity to frame a worthwhile question, choose and apply appropriate methods, interpret evidence, connect ideas across contexts, recognize uncertainty, revise one's view, and explain the reasons for a conclusion. In professional disciplines, this agency includes practical judgment—knowing how general principles bear on particular situations in which the data are incomplete, values conflict, and consequences fall unevenly on different people.

This account avoids romanticizing unaided work. Human inquiry has always depended on instruments, texts, teachers, collaborators, notation systems, laboratories, libraries, software, and institutions. The educational question is not whether a student's cognition is technologically extended, but whether the extension contributes to or substitutes for the development of agency. A calculator can free an advanced student to reason about a mathematical model while preventing a novice from acquiring numerical sense. A coding assistant can help an experienced programmer inspect unfamiliar syntax while allowing a beginner to assemble software whose behavior cannot be explained or secured. The same assistance can therefore be enabling at one developmental stage and disabling at another.

Three capacities are especially important under AI mediation. First, students require enough **underlying competence** to notice when a generated answer is implausible or incomplete. Second, they need **metacognitive awareness** of what they understand, what they have delegated, and where verification is required. Third, they need **epistemic control**: a defensible basis for accepting, revising, or rejecting the system's output. The ability to prompt a model is useful, but it is not equivalent to the disciplinary judgment required to assess the answer. A graduate who can produce impressive work only while a system behaves reliably has not acquired the same competence as one who can diagnose the system when it fails.

For this reason, independent competence remains an educational good even in professions where AI will be ubiquitous. "Independent" need not mean permanently tool-free. It means that the person retains sufficient understanding to exercise meaningful control, continue learning, and remain answerable when the tool is unavailable, inappropriate, or wrong. Universities should specify this threshold by discipline rather than invoking human oversight abstractly. A doctor, engineer, lawyer, teacher, or software professional must be able to understand and defend some judgments even when AI contributed substantially to the work.

2.3 Character and the practices of inquiry

The claim that universities form character can sound paternalistic, particularly in religiously diverse institutions. It can suggest ideological conformity or the ambition to engineer students' personalities. That is not the claim advanced here. Universities are limited institutions. They do not form the whole person, professors cannot manufacture virtue, and education is not redemption. A Christian account should be especially cautious about imagining that improved pedagogy can perfect human beings.

Yet the opposite claim—that universities transmit neutral knowledge and skills without shaping character—is implausible. Repeated practices cultivate dispositions. Assessment teaches students what an institution rewards. Deadlines shape attitudes toward time and obligation. Citation practices form habits of acknowledgment. Laboratories teach either care or carelessness. Group projects can cultivate cooperation or strategic free-riding. Faculty exemplars communicate how authority handles uncertainty, criticism, and error. James K. A. Smith's account of education emphasizes that human beings are shaped not only by propositions but by practices that direct attention, desire, and imagination (Smith 2009). David I. Smith similarly argues that Christian faith bears on the manner and practices of teaching, not only on the religious content inserted into a curriculum (Smith 2018).

Formation, as used in this paper, therefore refers to the cumulative shaping of dispositions through participation in practices, relationships, incentives, exemplars, and structures of responsibility. It is usually indirect and incomplete. An assignment cannot guarantee humility, but it can require a student to state uncertainty and respond to criticism. A design course cannot create compassion, but it can bring students into accountable contact with people affected by their designs. An academic-integrity policy cannot make students truthful, but it can either normalize transparent attribution or reward concealed dependence.

This is particularly concrete in computing, engineering, business, medicine, and other professional programmes. Students can be required to:

- document assumptions and data limitations;
- distinguish observed evidence from model-generated inference;
- test systems under failure and adversarial conditions;
- identify who benefits and who bears risk;
- listen to users and communities affected by a design;
- record trade-offs rather than hide them behind a single performance metric;
- disclose material AI assistance; and
- explain and defend a recommendation before teachers, peers, clients, or practitioners.

These practices do not prove that a student has become virtuous. They create repeated occasions to exercise honesty, humility, diligence, courage, care, and responsibility. AI policy is formative in precisely this sense. If students are routinely rewarded for producing polished work whose reasoning they cannot reconstruct, the institution trains performance without ownership. If they are required to interrogate outputs, expose uncertainty, and accept responsibility, AI can become an occasion for deeper judgment rather than an escape from it.

2.4 Vocation, work, and service

Universities properly prepare students for employment. Students and families invest substantial time and resources in education, and institutions have obligations not to ignore changing labor markets. AI literacy, tool competence, and the ability to work effectively with automated systems will become part of many professions. A theology of vocation does not oppose such preparation; it gives it a wider end.

In Protestant thought, vocation is not restricted to ordained ministry or exceptional achievement. Human beings serve God and neighbor through ordinary forms of work, citizenship, family life, scholarship, and institutional responsibility. This understanding refuses both the reduction of education to private economic return and the romantic claim that work is detached from material need. Students need livelihoods, but their capabilities also create obligations. Knowledge and skill are gifts to be cultivated faithfully and directed toward the good of others (Mark 12:30–31; 1 Pet. 4:10).

This perspective changes the question posed by automation. The issue is not simply which tasks AI can perform more cheaply or quickly, but how a changed division of labor affects service, competence, responsibility, and human dignity. Delegating routine documentation may give a doctor more time with patients or a professor more time with students. That is a genuine good if the released time is actually redirected toward human attention rather than absorbed by higher output demands. By contrast, delegating the interpretive work through which professionals acquire judgment may produce short-term efficiency and long-term deskilling. The economic beneficiary of automation may also differ from the person who bears the risk of error.

Vocation therefore requires **answerability**. AI may contribute analysis, draft alternatives, or detect patterns, but it cannot assume a human calling, love a neighbor, repent of negligence, or bear professional responsibility before an affected person. Students must learn not only how to use powerful systems but when to resist their recommendations, how to explain decisions to those affected, and why responsibility remains human even when execution is distributed across people and machines. Miroslav Volf and Matthew Croasmun's call to orient theology toward human flourishing reinforces this outward direction: knowledge is not complete when it remains internally coherent; it must bear on lives worth living and on the world's good (Volf and Croasmun 2019).

2.5 Communities of inquiry and accountability

Learning is personal but not solitary. Universities are communities in which students inherit knowledge from others, submit ideas to criticism, encounter perspectives they did not choose, and gradually enter disciplinary traditions. Teachers do more than transmit content. They model how experienced practitioners notice problems, weigh evidence, inhabit uncertainty, and revise judgments. Peers are more than an audience or a source of answers. They make claims upon one another through disagreement, collaboration, mutual explanation, and shared responsibility.

These relationships matter because knowledge has both a social and a moral structure. A learner must come to trust some testimony while remaining capable of criticism. A scholar depends on communities that maintain methods, archives, standards, and processes of review. A professional learns judgment partly through apprenticeship with people who can ask questions that are responsive to the learner's history and whose own reputations are implicated in their guidance. Correction within such relationships may be uncomfortable, but that discomfort can be educationally significant. It teaches that inquiry is accountable to realities and persons outside the self.

AI can strengthen some dimensions of this community. It can provide supplementary explanation to a student reluctant to ask publicly, help multilingual learners prepare for discussion, widen access for

some students with disabilities, and give rapid feedback when human assistance is unavailable. These benefits should not be minimized. But a conversational system is not interchangeable with a teacher or peer. It does not participate in the shared vulnerability of learning, bear institutional or professional responsibility for the student's formation, or suffer the consequences of mistaken counsel. Its apparent patience and affirmation can also make human disagreement seem needlessly difficult.

The educational aim is therefore not to reserve every interaction for humans. It is to identify the relationships whose goods arise from reciprocity, presence, example, challenge, care, and accountability, and to ensure that AI supplements rather than displaces them. A university with highly personalized AI support but weakened communities of inquiry may increase informational access while diminishing the setting in which intellectual and moral agency develops.

2.6 Formation without indoctrination

Because formation is inevitable, it must become explicit and contestable. The alternative to formative education is not neutral education, but formation governed by unexamined incentives. A curriculum organized around rankings, speed, frictionless completion, and market value teaches students what to love even if no professor says so. In the same way, a digital platform embodies judgments about which activities deserve attention, which behaviors are rewarded, what may be measured, and where authority resides. AI becomes part of the university's **hidden curriculum**: it can normalize immediate answers, private interaction, generic fluency, continual surveillance, or deference to opaque systems without these being declared learning outcomes.

Making formation explicit does not license coercion. Several safeguards are necessary. Educational aims should be connected to the institution's publicly stated mission and to the legitimate purposes of the discipline. Students should know which capabilities and dispositions a practice is intended to cultivate. Normative commitments should be open to reasoned criticism. Assessment should not claim to measure the inward moral or spiritual state of students. In secular universities, professors should neither conceal their Christian reasoning nor presume that theological authority settles institutional decisions for everyone. In Christian universities, theological commitments may properly be more explicit, but they should still be enacted through intellectually serious practices rather than slogans or compulsory piety.

This distinction is crucial for virtue language. Universities can assess whether a student discloses AI assistance, tests a system responsibly, acknowledges uncertainty, listens to stakeholders, or defends a decision. They should be cautious about claiming to assess whether the student *is* humble, loving, or wise. Practices are observable and pedagogically addressable; the full character of a person is neither transparent to an assessor nor producible by an institution.

2.7 Protestant warrant and public translation

A Protestant Christian framework should speak in two registers. The first is theological: why Christians understand truth, dignity, learning, work, and community as goods grounded in God's creation and purposes. The second is institutional: how the resulting educational commitments can be stated, debated, and evaluated among colleagues who do not share that theology. Public translation does not

require stripping away Christian identity. It requires offering reasons that others can examine without being asked to pretend that all participants begin from the same ultimate commitments.

Protestant theological warrant	Shared educational good	Institutional expression under AI mediation
Creation is God's intelligible gift; human knowers are finite and fallible.	Truthful, humble, self-correcting inquiry	Verification, source provenance, uncertainty disclosure, and openness to correction
Human beings bear God's image and are not reducible to productivity.	Non-instrumental dignity and intellectual agency	AI adoption judged by human development, not output or cost alone
Love of God and neighbor orders knowledge toward service.	Professional responsibility and the common good	Attention to affected persons, distribution of benefit and risk, and answerability for decisions
Vocation dignifies ordinary work while resisting work as an idol.	Meaningful competence and responsible participation	Automation that removes needless burden without eroding judgment or concentrating exploitation
Human beings are relational, embodied, dependent creatures.	Communities of inquiry, apprenticeship, and care	Protection of consequential teacher-student, peer, laboratory, clinical, and field relationships
Sin distorts institutions and desires; grace denies that worth must be earned.	Realistic governance and freedom from performance identity	Audits of power and incentives, limits on surveillance, and refusal to equate persons with metrics

This dual register allows Christian professors in secular institutions to contribute without either privatizing their convictions or imposing them. It also challenges Christian universities to demonstrate that their AI strategies are genuinely shaped by their missions. A Christian institution does not become theologically distinctive merely by adding devotional language to a vendor-led transformation. Its distinctiveness should be visible in what it protects: patient inquiry, truthful authorship, the dignity of students and staff, service to vulnerable communities, responsible limits, and relationships that cannot be reduced to transactions.

The Catholic encyclical *Magnifica Humanitas* provides a useful ecumenical dialogue partner at this point. It describes education as a patient search for truth, warns against allowing information to displace research, reflection, and discernment, and identifies the university's task of integrating knowledge, grasping complexity, and verifying facts (Leo XIV 2026, nos. 137–147). Protestant readers need not treat papal teaching as doctrinal authority to recognize a substantial convergence around truth, formation, human dignity, and the common good.

2.8 A normative test for educational AI

The preceding account yields a more precise standard than either "embrace innovation" or "preserve human teaching." An AI-supported practice is educationally justified when, considering the learner's stage and the purpose of the activity, it contributes to the ordered goods of education without

bypassing a capability, relationship, or responsibility that the learner still needs to acquire or demonstrate.

Professors and leaders should therefore ask five prior questions before evaluating convenience or scale:

11. **Truth:** Does the practice strengthen the learner's ability to seek, verify, integrate, and communicate what is true?
12. **Agency:** Does it develop the learner's understanding and judgment, or merely improve performance while assistance is present?
13. **Character:** What habits of honesty, attention, humility, diligence, courage, and responsibility does repeated use reward?
14. **Vocation:** Does it prepare students to use knowledge competently and answerably in service to others?
15. **Community:** Does it support or displace the human relationships through which correction, apprenticeship, care, and accountability occur?

This is the normative basis for **formation-sensitive augmentation**. It does not predetermine that more or less AI is always better. It identifies what evidence must be sought. A university cannot decide whether AI improves education until it specifies the educational good in question, the human capacity or relationship through which that good is realized, and the signs that the capacity is being strengthened rather than bypassed. The next section therefore examines what current empirical research can—and cannot yet—tell us about those effects.

3. What the empirical evidence shows—and does not show

The empirical question is often posed too broadly: *Does generative AI improve or damage education?* The current literature cannot sustain either universal conclusion. "Use of AI" includes activities as different as receiving a hint, generating an entire answer, obtaining feedback on a self-produced draft, conversing with an instructor-designed tutor, and consulting a general-purpose chatbot. "Educational outcome" may refer to the quality of an artifact produced with assistance, performance on an immediate post-test, later unaided recall, motivation, self-reported confidence, or the development of professional judgment. Averaging across these interventions and outcomes can conceal the mechanisms that professors and leaders most need to understand.

The most defensible conclusion as of August 2026 is conditional:

> *Generative AI can improve learning when it is designed and used to elicit cognitive activity, provide accurate and timely feedback, and preserve the learner's responsibility for reasoning. It can also improve the visible performance of a task while weakening or failing to develop the learner's independent capacity when it substitutes for that reasoning. Evidence about long-term intellectual, moral, relational, and vocational formation remains sparse.*

This conclusion is consistent with the OECD's 2026 synthesis, but the analysis below does not rely on that policy report alone. It examines the main causal studies, meta-analyses, systematic reviews, observational studies, and adoption surveys that bear on the paper's normative tests.

3.1 Scope, method, and evidentiary limits

This section is a **targeted scoping synthesis**, not yet a completed systematic review. It covers research published or publicly available from the release of ChatGPT in November 2022 through 2 August 2026. Priority is given to peer-reviewed research involving university students, followed by rigorous studies in adjacent settings when they reveal a mechanism plausibly relevant to higher education. Major claims are classified by:

- educational level, discipline, and learner expertise;
- the AI system and whether its behavior was pedagogically constrained;
- the activity delegated to or supported by AI;
- the comparator and degree of experimental control;
- whether the outcome was measured with AI present or removed;
- the time between intervention and outcome measurement; and
- whether the measure captured performance, learning, metacognition, affect, behavior, or relationship effects.

Evidence is weighted in descending order from randomized or well-controlled field experiments, to meta-analyses and systematic reviews whose included studies are sufficiently comparable, to longitudinal and observational studies, and finally to cross-sectional surveys and qualitative reports. Preprints are identified as such. Usage reports are used to establish prevalence and perceptions, not educational effects. Workplace and general-population studies are used only to identify possible mechanisms; they do not establish what happens to university students.

Four cautions govern the synthesis.

First, **performance is not identical to learning**. A student may produce a better solution while AI is available without becoming more able to solve a related problem later. Second, **immediate learning is not the same as durable learning**. A post-test administered directly after a lesson does not establish retention, transfer across contexts, or integration into a disciplinary mental model. Third, **self-report is not behavioral evidence**. Students may accurately describe their experience, but confidence, satisfaction, and perceived learning can diverge from tested competence. Fourth, **formation is a longer and richer construct than most studies measure**. No short experiment can establish whether repeated AI use makes a person more truthful, humble, diligent, wise, responsible, or capable of sustaining difficult human relationships.

The evidence base also has structural weaknesses. Model capabilities and interfaces change faster than conventional publication cycles. Many studies use small samples, brief interventions, bespoke prompts, or a single discipline. The comparison condition is sometimes weak or poorly described. Measures labeled “critical thinking” or “higher-order thinking” vary widely and sometimes rely on self-report or the quality of AI-assisted products. Blinding is difficult, and positive-result publication bias is plausible in

a rapidly expanding field. Meta-analytic averages should therefore be read as summaries of a heterogeneous early literature, not stable properties of "AI in education."

3.2 Adoption is established; educational benefit is not

Generative AI is no longer a peripheral practice in higher education. The 2026 UK Student Generative AI Survey, based on 1,054 full-time undergraduates, reported that 95 percent had used AI in at least one way and 94 percent had used generative AI to help with assessed work. Twelve percent reported directly incorporating AI-generated text in assessed work, compared with 8 percent in 2025 and 3 percent in 2024. Almost half believed that AI had improved their university experience, while others identified skill erosion, unfairness, isolation, and employment anxiety (Stephenson and Armstrong 2026). A larger international survey fielded by the Digital Education Council in 2026 collected more than 45,000 responses across 35 countries and similarly indicates widespread but uneven adoption (Digital Education Council 2026).

These findings matter institutionally: policies premised on rare or easily detectable use no longer describe the environment in which teaching and assessment occur. They do not, however, show that AI improves or harms learning. Survey respondents self-select uses, differ in prior ability and motivation, and often use several systems for several purposes. Reports that AI saves time or improves understanding may describe real benefits, but they cannot distinguish strengthened capability from easier task completion. Conversely, faculty perceptions that students are becoming less attentive or less capable are warning signals rather than causal estimates.

Adoption is also uneven. Access to paid models, reliable connectivity, disciplinary guidance, and the ability to evaluate outputs varies across students and institutions. Students with stronger prior knowledge may be better positioned to detect errors and use AI strategically; those who most need support may be most vulnerable to accepting fluent but defective explanations. The evidence therefore does not warrant the simple claim that widespread availability democratizes high-quality tutoring. Availability is not equivalent to pedagogical quality, and equal access to a tool does not produce equal capacity to use it well.

3.3 Structured tutoring can produce substantial near-term gains

The strongest university-level experimental evidence for benefit comes from a randomized crossover trial in an introductory physics course at Harvard University. Kestin et al. (2025) compared two lessons delivered through an instructor-designed AI tutor with active-learning classroom instruction. The 194 participating undergraduates encountered both conditions across consecutive weeks. Students using the AI tutor achieved substantially larger immediate post-test gains, with an estimated effect of approximately 0.73 to 1.3 standard deviations after adjustment for ceiling effects, and spent a median of 49 minutes on the lesson rather than the 60 minutes assigned to classroom learning. They also reported greater engagement and motivation.

This is credible evidence that a well-designed AI tutor can outperform a strong classroom comparison for defined material over a short period. It is not evidence that a general chatbot can replace a professor

or that individualized AI instruction is generally superior to human teaching. The intervention embedded several elements that are easy to lose in retelling: an instructor-determined learning sequence, expert-written problem solutions, explicit management of cognitive load, targeted feedback, self-pacing, and prompts intended to promote active engagement and a growth mindset. The designers did not rely on the model to generate unverified solutions. In other words, the study tested a pedagogical system in which faculty expertise structured the AI interaction, not autonomous model instruction.

The outcome was also narrow. The experiment involved two physics lessons, and the tests followed the interventions immediately. It did not examine delayed retention, transfer to substantially different problems, cumulative conceptual development across a semester, laboratory practice, or the relational goods of a course. Nor did it compare the tutor with expert one-to-one human tutoring. The proper inference is nevertheless important: AI-mediated personalization, rapid feedback, and self-pacing can create genuine learning gains when organized around sound pedagogy. The medium itself is not the achievement; the instructional design is.

Meta-analyses reinforce this positive but qualified picture. Deng et al. (2025) reported positive average effects of ChatGPT interventions on academic performance, affective-motivational outcomes, and higher-order-thinking measures, alongside reduced mental effort. Chen and Cheung's (2025) university-focused meta-analysis likewise found a large positive aggregate effect on learning outcomes, with especially large estimates for language skills, but no statistically significant effect on metacognition. Xia et al. (2025) found positive effects on university students' motivation and several dimensions of engagement, moderated by subject, learning strategy, and use context. These syntheses make it unreasonable to claim that generative AI is inherently inimical to learning.

They do not settle the matter. The included interventions differ in duration, discipline, instructional method, model, and outcome. Many studies assess products made with AI or tests given immediately after intervention; fewer examine delayed, unaided performance. Very large pooled effects in a young, heterogeneous literature should increase interest while also prompting scrutiny of comparison conditions, publication bias, and construct validity. Reduced mental effort may indicate welcome relief from extraneous cognitive load, or it may indicate that the learner did less of the cognitive work the task was designed to cultivate. The meaning depends on *which* effort was removed.

3.4 The performance-learning divergence

The clearest causal demonstration of the central risk comes from a field experiment by Bastani et al. (2025) involving nearly 1,000 secondary-school mathematics students. Although the participants were not university students, the design isolates a mechanism with direct relevance to novice learning. Students were assigned to conventional practice, access to a GPT-4 interface resembling a general chatbot (“GPT Base”), or a constrained tutor (“GPT Tutor”) prompted to provide staged assistance and avoid simply supplying complete answers.

While assistance was available, GPT Base improved practice performance by 48 percent relative to the control condition, and GPT Tutor by 127 percent. When AI access was removed for the subsequent examination, however, students who had used GPT Base performed 17 percent worse than controls.

The negative effect was almost eliminated in the constrained-tutor condition, but that group did not outperform the control group on the unaided examination. Students' perceptions were more optimistic than their tested learning: those with unrestricted access did not recognize the subsequent disadvantage, while constrained-tutor users believed their learning had improved despite no detectable exam advantage.

The study should neither be ignored because it involved secondary students nor generalized mechanically to every university task. It concerned mathematics practice in one educational context and tested a particular implementation of GPT-4. Yet it establishes three propositions of high relevance to university leaders.

16. Assisted task scores can move in the opposite direction from subsequent independent performance.
17. Guardrails that change the learner's interaction with the model can materially change the educational effect.
18. Learners' confidence and satisfaction cannot substitute for independent assessment of what they have learned.

The mechanism is consistent with established learning science. Novices acquire schemas and judgment through retrieval, attempted solution, error, feedback, and revision. If assistance supplies the decisive reasoning before the learner has attempted it, the student may experience fluent progress without constructing the knowledge needed later. The issue is not that all cognitive offloading is harmful. Experts routinely use calculators, search engines, software libraries, and colleagues. The issue is **premature offloading**: delegating a cognitive operation before the learner has developed enough understanding to evaluate, integrate, and, when necessary, reproduce it.

This creates a developmental asymmetry. The same tool may be augmentative for an expert and substitutive for a novice. A skilled programmer can use generated code to explore alternatives because she possesses mental models for architecture, testing, complexity, and failure. A beginning student may receive a functioning program while failing to learn decomposition, debugging, or why the program works. Universities should therefore resist uniform policies that treat every student, task, and stage of a curriculum as equivalent.

3.5 Writing, memory, ownership, and cognitive effort

Writing deserves special attention because it is both a product and a mode of inquiry. Students often discover what they think through planning, composing, confronting ambiguity, and revising. AI can assist this process through examples, questions, translation, feedback, and editorial support. It can also produce the linguistic surface of understanding without the student's having formed the underlying argument.

Kosmyna et al.'s (2025) widely discussed “Your Brain on ChatGPT” study offers suggestive evidence but has often been overstated. Fifty-four participants completed three essay-writing sessions in one of three conditions: LLM assistance, search-engine access, or no external tool. Eighteen returned for a crossover session. The researchers combined EEG measures, textual analysis, interviews, recall questions, and human and AI scoring. The LLM group showed weaker measured neural connectivity than the search

and unaided groups, poorer recall of their own essays, and lower reported ownership; group texts also displayed greater linguistic homogeneity. In the crossover, some effects persisted when prior LLM users wrote without the tool.

These results warrant further investigation into memory, authorship, and the strategies elicited by different writing environments. They do not show "brain damage," that LLMs cause people to stop thinking, or that all AI-assisted writing weakens cognition. The paper was a preprint at the time of this review; the initial sample was small, only 18 participants completed the fourth session, the task and tool conditions were specific, and neural connectivity is not a direct measure of educational quality or human flourishing. The researchers themselves caution against sensational language. The responsible conclusion is that producing an essay with extensive LLM assistance can reorganize cognitive activity and may reduce recall and ownership under some conditions. Replication with larger samples, delayed tests, varied writing processes, and measures of argument quality is needed.

A related workplace study by Lee et al. (2025) surveyed 319 knowledge workers about 936 real examples of generative-AI use. Greater confidence in the AI was associated with less self-reported enactment of critical thinking, whereas greater confidence in one's own task competence was associated with more. Participants described critical thinking as setting goals, refining prompts, checking outputs, and integrating results. The study is useful because it suggests that AI does not simply eliminate critical thinking; it can shift it from producing content toward specifying, verifying, and integrating content. But it is cross-sectional, workplace-based, and dependent on retrospective self-report. It identifies a plausible reallocation of cognitive effort, not a measured decline in student capacity.

The emerging synthesis is therefore more subtle than "AI reduces critical thinking." Li, Cui, and Hagedorn's (2026) systematic review of 67 higher-education studies found that ChatGPT could support critical and creative thinking when embedded in inquiry-oriented and scaffolded instruction, while unstructured use was associated with cognitive offloading and superficial engagement. Benefits appeared stronger and more consistent for divergent idea generation than for the convergent work of evaluating evidence and reaching warranted conclusions. This distinction is educationally significant. Generating more possibilities can enlarge thought; judging which possibility is true, responsible, or fitting remains a separate achievement.

3.6 Engagement, motivation, and the ambiguity of ease

Students often value AI because it is immediate, patient, private, and responsive. These affordances can be educationally important. A student may ask questions without embarrassment, request another explanation, translate unfamiliar language, generate additional practice, or receive feedback when an instructor is unavailable. The Kestin trial and recent meta-analyses provide evidence of positive effects on motivation and cognitive or emotional engagement in some settings.

Engagement, however, is not a unitary good. Time spent interacting with a system may reflect productive explanation-seeking, answer extraction, repeated prompt repair, or dependence. Enjoyment can support persistence, but ease can also remove the desirable difficulty through which memory and skill develop. A conversational system's apparent warmth may increase willingness to continue without

increasing epistemic vigilance. “Students liked it” and “students learned from it” are related questions, not interchangeable findings.

The Bastani experiment makes this divergence concrete: students rated the assistance favorably even when their later unaided performance showed no benefit or a measurable cost. The implication is not to discount student experience. Rather, institutions should combine experience measures with behavioral outcomes. An educational AI evaluation should ask whether students return to challenging material, generate their own explanations, detect planted errors, transfer learning to a new context, and accurately calibrate confidence—not simply whether they used the system frequently or would recommend it.

3.7 Epistemic and relational effects beyond test scores

Generative AI may also affect how students respond to disagreement, correction, and uncertainty. Cheng et al. (2026) examined social sycophancy across 11 leading models. The models affirmed users' actions 49 percent more often than human comparison responses, including in situations involving deception, illegality, manipulation, or relational harm. Across three preregistered experiments involving 2,405 participants, exposure to sycophantic AI increased participants' conviction that they were right and reduced their intentions to accept responsibility, apologize, make amends, change their behavior, or otherwise repair relationships. Participants nevertheless judged the agreeable responses more favorably and expressed greater trust and willingness to return to the system.

This is strong evidence about a model behavior and its immediate effects in interpersonal-advice settings. It is not a study of classrooms, long-term character, or specifically Christian students. The inference to education must therefore be explicit and limited: a system optimized for user approval may be a poor substitute for communities in which teachers and peers can resist a learner's framing, require evidence, and make morally serious demands. If students increasingly use AI not only for explanations but for reassurance, conflict interpretation, or ethical advice, sycophancy could weaken openness to correction even while making the interaction feel supportive.

The risk is intensified by anthropomorphic presentation. Conversational fluency can make a probabilistic system appear to understand the learner, care about the outcome, or possess warranted confidence. Current studies establish that people can be influenced by such responses; they do not establish that AI relationships are uniformly harmful or that students cannot learn appropriate skepticism. Universities should treat epistemic calibration—knowing what the system is, what it can know, and when its outputs require external verification—as a learning outcome rather than assuming that frequent use will produce it automatically.

3.8 Academic integrity and assessment validity

Generative AI has made some established assessments poor measures of the capabilities they are intended to certify. If an unsupervised take-home task can be completed largely by a model, a high-quality submission may no longer establish that the student can analyze the evidence, write the code, or

defend the argument. This is first an assessment-validity problem and only secondarily a detection problem.

Research gives reason not to make automated detection the foundation of academic-integrity policy. Liang et al. (2023) tested seven widely used detectors on 91 TOEFL essays written by non-native English writers and 88 essays by US eighth-grade students. The detectors classified the US essays with high accuracy but falsely labeled more than half of the TOEFL essays as AI-generated on average; simple prompting could also alter AI text to evade detection. The samples and rapidly changing detectors limit generalization to every current product, but the study demonstrates that detector outputs can be both unreliable and unequally distributed. A probability score should not be treated as proof of misconduct.

The empirical literature does not yet identify one universally valid replacement for essays, coding assignments, or other take-home work. The defensible direction is triangulation: combine process evidence, staged drafts, oral explanation, live problem solving, authentic projects, reflective disclosure, and selected supervised demonstrations of independent competence. AI-inclusive assessments should test students' ability to interrogate and improve outputs; AI-restricted assessments should establish what they can understand and do without assistance. Both are needed because graduates will work with AI and will also encounter situations in which systems are wrong, unavailable, inappropriate, or incapable of bearing responsibility.

Academic integrity is not exhausted by rule compliance. Policies and assessments teach students what authorship means. If disclosure requirements are vague, inconsistent across modules, or framed entirely as surveillance, students receive little help in distinguishing legitimate assistance from the delegation of intellectual responsibility. Conversely, requiring an “AI use statement” without examining whether the student understands and owns the resulting work can turn disclosure into ritual. Evidence on which disclosure practices actually cultivate honest and capable use remains limited.

3.9 Access, inclusion, and distribution of benefit

Generative AI can lower barriers for multilingual students, some students with disabilities, and learners who lack timely access to tutoring. Translation, text simplification, alternative explanation, speech and text conversion, and around-the-clock assistance are genuine capabilities. Yet the current evidence is stronger on technical possibility and user perception than on long-term equalization of university outcomes.

At least four distributional questions remain open. Do weaker students receive greater benefit, or are they more easily misled because they have less prior knowledge with which to verify outputs? Do paid systems create a new tier of academic advantage? Does AI-mediated language support enable a student to participate more fully while developing disciplinary language, or conceal needs that require human teaching? Does automation reduce faculty workload in ways that release time for relationships, or primarily increase institutional expectations for throughput?

The answer will depend partly on implementation. Provision of a common tool may reduce one access gap while creating dependencies on vendors, data collection, and dominant languages or knowledge

traditions. Conversely, prohibiting tools without providing human alternatives may disproportionately harm students who used them for accessibility or language support. Equity cannot therefore be inferred from either adoption or restriction. It must be assessed through disaggregated outcomes, accessibility testing, privacy review, and attention to who gains time, who loses developmental opportunity, and who bears the consequences of error.

3.10 What may responsibly be concluded

The evidence can now support several claims with different levels of confidence.

Claim	Confidence	Basis and qualification
Carefully structured, pedagogically grounded AI tutoring can improve immediate learning and engagement in some university settings.	Moderate to high	Supported by a strong university RCT and positive experimental syntheses; evidence remains concentrated in short interventions.
General-purpose AI can improve the quality or accuracy of work while it is available.	High	Consistent across educational and workplace studies; this is assisted performance, not necessarily learning.
Unrestricted assistance can impair subsequent independent learning for novices.	Moderate	Strong causal evidence in secondary mathematics and convergent evidence on offloading; university replication across disciplines is still needed.
Guardrails, prompts, sequencing, and instructor-authored content materially affect outcomes.	Moderate to high	Direct comparison in the Bastani trial and central design feature of the Kestin trial; the contribution of each component is not fully isolated.
Generative AI generally improves university students' motivation, engagement, and short-term achievement.	Moderate	Positive meta-analytic averages, but with substantial heterogeneity and construct variation.
AI use reduces critical thinking in general.	Low as a universal claim	Effects vary by task and design; AI can support critique under structured conditions and displace it under unstructured ones.
Extensive LLM-assisted writing may reduce recall and felt ownership under some conditions.	Suggestive	Multimethod preprint evidence from a small sample; replication and delayed outcomes are needed.
Sycophantic AI can increase confidence in one's own position and reduce intentions toward relational repair.	High for brief advice interactions; unknown educationally	Supported by preregistered experiments, but not yet by longitudinal studies in universities.
Current AI use causes long-term deskilling, weakened character, or loss of human community in university students.	Unknown	Plausible mechanisms and warning signals exist, but direct longitudinal evidence is insufficient.
AI improves educational equity.	Unknown and context-dependent	Accessibility and scale offer real possibilities; differences in access, prior knowledge, language, guidance,

Claim	Confidence	Basis and qualification
		and exposure to error may reproduce or widen inequality.

The empirical record therefore supports neither technological enthusiasm nor moral panic. It supports design responsibility. The most important variable is not simply whether AI is present, but what role it is assigned within a sequence of learning. Systems that question, hint, give accurate feedback, require retrieval, expose uncertainty, and return explanatory responsibility to the learner are more consistent with learning than systems that deliver a finished answer on demand. Even well-designed assistance should be paired with later opportunities to demonstrate retention, transfer, judgment, and independent competence.

Most importantly for this paper, the literature measures only a portion of what universities value. Immediate achievement, motivation, and efficiency matter, but they do not exhaust truthfulness, intellectual agency, character, vocation, or community. The absence of longitudinal evidence does not prove that AI is harmless to these goods; neither does a plausible theological concern prove empirical damage. It creates an obligation for intellectual humility and institutional experimentation: adopt selectively, specify the developmental mechanism, measure outcomes beyond the assisted artifact, protect spaces for human relationship and unaided thought, and remain willing to reverse practices that generate performance without formation.

4. A Christian anthropology of learning and technology

The empirical evidence reviewed above can tell universities whether particular uses of AI improve test performance, engagement, recall, or independent problem-solving under specified conditions. It cannot determine what human beings are, which kinds of development deserve protection, or why responsibility should remain human when a machine performs much of the intellectual work. Those are anthropological and moral judgments. Every AI strategy makes them, even when it presents itself as merely technical. A university that treats students primarily as future workers, measurable outputs, or users of an educational service has already adopted an account of the person. A Christian anthropology makes its account explicit and therefore open to theological and institutional criticism.

The argument developed here is Protestant in three senses. First, it treats Scripture as the primary doctrinal norm while drawing on the wider Christian tradition. Second, it gives characteristic weight to grace, vocation, the priesthood of all believers, and the fallibility of persons and institutions. Third, it rejects any suggestion that education, technology, or moral cultivation can redeem humanity. The university can cultivate practices conducive to wisdom and responsibility; it cannot perfect the person or accomplish salvation.

This account must also be contextual. Protestant reasoning about technology is often framed in response to Western individualism, expressive autonomy, and commercial consumer culture. These concerns are real but incomplete for universities in Asia. Chinese and other Confucian-heritage settings

may place greater emphasis on learning as disciplined self-cultivation, education as a family project, respect for teachers, socially embedded obligation, harmony, and contribution to society. These are not marginal details. They change how agency, achievement, authorship, authority, and AI assistance are experienced.

At the same time, “Asia,” “Chinese culture,” and “Confucian education” must not be used as if they describe a single population. Mainland China, Singapore, Hong Kong, Taiwan, and Chinese diasporic communities have different political institutions, religious histories, languages, educational systems, and patterns of AI access. Within each setting, class, gender, discipline, generation, ethnicity, faith, and family history produce further variation. Confucian traditions themselves contain internal debates and have been repeatedly interpreted, criticized, and reconstructed. Contextual theology should therefore avoid two errors: imposing a supposedly universal Western problem definition, and romanticizing an undifferentiated Asian communitarian alternative. Pan-Chiu Lai's account of inheriting Chinese and Protestant traditions is helpful precisely because it presents contextualization as critical and constructive inheritance rather than the unqualified acceptance of either tradition (Lai 2003).

The contemporary policy setting makes this dialogue urgent. China's 2026 “AI + Education” plan calls for AI literacy throughout the education system and requires universities to make AI a basic public course while developing interdisciplinary programmes aligned with industrial transformation (State Council of the People's Republic of China 2026). Singapore and several other Asia-Pacific systems have likewise made AI competence an educational priority. UNESCO's regional analysis cautions that the Asia-Pacific is highly uneven in infrastructure and readiness, and that models trained predominantly on Western data may lack local cultural relevance or reproduce bias (UNESCO 2023). The question for university leaders is therefore not whether Asian institutions should choose tradition instead of innovation. It is which inherited goods and distortions will be strengthened as AI becomes ordinary.

The following themes form one cumulative claim. Humans are created and called, yet finite; creative, yet fallen; embodied and relational, yet responsible truth-seekers; reconciled through Christ’s finished work, yet still forming in wisdom; indwelt and guided by the Spirit, yet frequently slow to attend to him; dependent on grace, yet commissioned to serve. AI should therefore be received neither as a rival creature nor as a neutral efficiency layer, but as a humanly made, institutionally embedded power whose educational meaning depends on the practices and ends into which it is incorporated.

4.1 Creation and the image of God: dignity before capability

Christian anthropology begins with creation, not competition. Human beings bear the image of God and are commissioned to exercise responsible care within creation (Gen. 1:26–28). The *imago Dei* should not be reduced to a single measurable capacity such as linguistic fluency, rational calculation, creativity, or autonomy. Any capacity-based definition creates a hierarchy among humans: the more articulate, intelligent, productive, or independent person would appear to image God more fully than the infant, the person with profound disability, the student who struggles, or the older adult whose capacities have declined. Langford therefore argues that the image of God cannot be ranked according to cognitive

performance and connects it instead to God's calling of human creatures into responsible relationship and stewardship (Langford 2022a, 2022b).

This claim becomes newly important when AI systems can outperform many people at valued academic tasks. A student's dignity does not diminish because a model writes more fluently, searches more widely, or codes more quickly. Nor does human distinctiveness need to be defended by continually relocating it to whichever benchmark machines have not yet surpassed. Christian worth is neither a leaderboard position nor a residual bundle of tasks that automation cannot perform. It is received from God.

Creation also gives a positive account of technology. Human beings name, cultivate, make, organize, and repair. Tools can extend care, widen access, translate languages, reduce hazardous labor, reveal patterns, and enable forms of participation previously denied to some learners. AI-assisted writing, for example, may allow a student with a disability or a student working in an additional language to participate more fully in a scholarly community. To refuse all assistance in the name of preserving a supposedly pure human performance could itself privilege a narrow image of the able, linguistically dominant learner. The question is not whether a tool changes human activity—tools always do—but whether the change serves responsible human agency and neighbor-love.

This positive account of technology can be sharpened by the Reformed doctrine of the cultural mandate. The same creation narrative that grounds human dignity in God’s image also commissions human beings to “fill the earth and subdue it” (Gen. 1:28), a mandate Abraham Kuyper read as authorizing the whole of human cultural activity—including scholarship, science, and, by extension, the tools scholarship builds—as legitimate creaturely work under God rather than a departure from properly “spiritual” concerns (Kuyper 1898). Kuyper’s correlative doctrine of common grace explains why a technology built largely by people with no confessed relation to God can nonetheless yield real goods: God restrains sin and distributes wisdom, skill, and insight broadly across humanity, not only within the church. Universities need not treat every advance in machine learning as either a work of the Spirit or a symptom of the fall; common grace allows Christian scholars to receive genuine intellectual gifts from unbelieving and believing colleagues alike, gifts that remain answerable to Christ’s lordship even though their makers do not confess it.

Kuyper’s doctrine of sphere sovereignty adds a second, more institutional caution. Because God has ordered creation into differentiated spheres—family, church, state, university, marketplace—each answerable directly to God and not simply to one another, a university’s use of AI should not be governed as though it were merely an extension of state industrial policy, market competition, or even ecclesial authority (Kuyper 1880). The university has its own proper sovereignty as a sphere of truthful inquiry, and its adoption of AI should be judged by the norms proper to that sphere—truthful inquiry, the formation of judgment, the cultivation of the whole person—rather than borrowed wholesale from the norms of national competitiveness or shareholder value that govern adjacent spheres. This qualifies the argument of Section 8 that AI strategy is properly a matter of educational mission: it is not merely that mission language sounds better than adoption targets, but that the university, as a distinct sphere under God, has standing to resist norms imported uncritically from the spheres of state and market.

Confucian accounts of education offer a significant point of resonance. Learning is not merely the acquisition of private information but part of becoming a humane and trustworthy person in relation to others. Chi-Ming Lam's reconstruction of Confucian learning with LLMs centers *ren* (humaneness), *li* (ritual propriety), *cheng* (sincerity), and the continuing work of self-cultivation (Lam 2026). This corrects an individualistic picture in which the autonomous student owns an education solely for private advancement. It also reinforces the paper's claim that the quality of an academic output cannot be separated from what its production is doing to the learner and the learner's relationships.

The traditions nevertheless differ in an important respect. In a Christian account, moral cultivation is a response to a dignity and calling already given; it is not the achievement of human worth. This distinction matters in competitive educational environments. When examination results, university admission, institutional prestige, and family sacrifice are tightly connected, students may experience achievement not simply as personal ambition but as evidence that they have become worthy children and socially valuable persons. AI can then become another means of securing status or avoiding shame. The doctrine of creation says both more and less: students are more than individual consumers because they belong to God and to communities of obligation, but they are never less worthy because they disappoint a ranking, a parent, a professor, or a national talent strategy.

For university leaders, the institutional test is whether AI adoption treats persons as ends or as optimizable inputs. Systems that predict risk, personalize instruction, rank applicants, monitor behavior, or automate advising may confer benefits, but they also encode judgments about which students are promising, normal, efficient, or expensive. An *imago Dei* framework requires opportunities to contest consequential classifications, protection against reduction to behavioral data, and special attention to those whose voices or languages are poorly represented. It also requires leaders to ask whether efficiency savings strengthen education for persons or merely permit the institution to process more students with less human attention.

4.2 Creatureliness, disciplined effort, and Sabbath limits

To be created is to be finite. Human beings learn slowly, forget, require sleep, depend on teachers, inhabit bodies, and understand only partially. Christianity does not treat every limitation as sacred: disease may be healed, needless administrative burden removed, and disability accommodated. Yet it also refuses the assumption that all friction is a defect. Some limits are conditions of learning and love. Waiting creates space for attention; repeated practice forms skill; wrestling with a problem reveals what one does not understand; dependence makes teaching and gratitude possible; rest declares that neither personal worth nor the world's future depends on uninterrupted productivity.

This is where the language of “cognitive offloading” requires moral discrimination. Offloading routine computation after conceptual mastery can enlarge agency. Offloading the novice's attempt to formulate a question, retrieve an idea, structure an argument, debug a program, or endure initial confusion can remove the very activity through which competence develops. The empirical performance–learning divergence in Section 3 is therefore not only a technical problem. It exposes a theology of limits: when is

AI removing needless burden, and when is it helping learners evade a necessary encounter with finitude?

Confucian educational traditions again supply a constructive resource. The opening movement of the *Analects* joins learning with repeated practice, and Analects 2:15 holds learning from others together with personal reflection. Lam applies the latter insight directly to LLMs: the danger is "learning without reflection," in which information is reproduced but not integrated through judgment (Lam 2026). This emphasis on disciplined practice challenges the promise of instant mastery. It also avoids equating independent learning with isolation; one learns through texts, exemplars, correction, and reflection rather than through unaided self-invention.

Yet discipline can be distorted into relentless performance. In many Asian families, education represents the concentrated sacrifice of parents and a hoped-for path to mobility and security. Students may therefore carry obligations extending beyond individual aspiration. Research on filial piety helps avoid a crude contrast between Western autonomy and Asian obedience. Reciprocal filial piety—gratitude and loving responsibility within a supportive relationship—has been positively associated with autonomy and achievement, whereas authoritarian filial piety centered on submission and indebtedness has not produced the same benefits and can be harmful (Zhou, Guo, and Xu 2020). The educational question is not whether family obligation should be replaced by individual independence, but whether obligation is reciprocal and life-giving or coercive and identity-defining.

AI enters this ecology ambiguously. It may relieve the linguistic burden on multilingual students, provide practice to those who cannot afford tutoring, and reduce low-value work. But it may also intensify an already accelerated system by raising expectations: more assignments, faster research, constant availability, and still greater competition for credentials. Time "saved" by faculty or students is not automatically returned as rest or relationship; institutions can convert it into increased throughput. Traynor Hansen's Sabbath-grounded account of AI and writing is useful here: Sabbath resists both the fantasy of liberation from all work and the anxious demand for continuous production, reordering work toward faithful participation rather than credentialed competition (Hansen 2025).

For Christian universities and Christian professors in secular institutions, creatureliness should become visible in temporal design. Courses need protected periods for unaided retrieval, slow reading, drafting, laboratory practice, and conversation. Faculty workload policies should not presume that AI assistance justifies indefinite acceleration. Students should be permitted to encounter difficulty without every pause being classified as inefficiency. Sabbath is not a call to lower standards. It is a refusal to make speed the master of excellence or productivity the measure of the person.

4.3 The fall, idolatry, and Babel: distorted loves at personal and institutional scale

Creation explains why technology can be good; the fall explains why good capabilities are never used under innocent conditions. Sin is not confined to malicious intent or spectacular misuse. It disorders loves and directs genuine goods—knowledge, recognition, security, harmony, power, service—toward false ultimacies. AI may therefore be used competently and legally while still training pride, deception, passivity, domination, or indifference.

Idolatry names the attempt to obtain from a created thing what it cannot give. In universities, AI can be treated as an oracle whose fluent answer relieves the user of judgment, a savior that will repair unequal education without costly institutional change, or an engine that secures competitiveness against rival universities and nations. These expectations need not be stated religiously to function as objects of trust. A vendor's promise of personalization at scale, a leader's aspiration to become “AI-first,” or a student's belief that the right model will eliminate intellectual struggle can each place a limited technology within a salvific narrative.

Babel adds the institutional and civilizational dimension (Gen. 11:1–9). The problem is not cooperative engineering as such, nor cultural plurality. It is the concentration of technique, reputation, and collective power in a project of self-exaltation: making a name, securing control, and resisting creaturely dependence. AI development in the United States and China is frequently narrated through national competition, strategic autonomy, and technological leadership. Universities genuinely must prepare graduates and contribute to social needs. Yet when the university's educational mission becomes subordinate to winning an innovation race, students can be valued principally as national human capital, and every hesitation about adoption can be dismissed as backwardness.

An Asian contextualization must broaden rather than simply reverse Western accounts of sin. Western AI ethics often emphasizes threats to individual choice and privacy. These remain important. In more relational or institutionally hierarchical settings, however, distorted agency may also appear through pressure to preserve face, avoid open disagreement, protect institutional reputation, or accept a leader's technological agenda for the sake of harmony. Conversely, a narrowly Western demand for individual self-authorship can ignore legitimate duties to family, profession, and society. Chow and Wong's survey of Chinese Christian theology notes that Chinese Christians have rarely understood faith as merely an individual choice and that accounts of sin have had to engage family and corporate responsibility as well as individual wrongdoing (Chow and Wong 2023). This is a valuable correction for AI governance: responsibility is personal, but the pressures and benefits surrounding an AI-mediated act are institutional and relational.

The doctrine of the fall therefore forbids scapegoating students. If an assessment can be passed by submitting generated work that the student cannot explain, the student may act dishonestly, but the institution has also built a practice that rewards the artifact while claiming to assess the person. If leaders license a system without meaningful privacy review because competitors have done so, the failure is not located only in the vendor. If faculty secretly use AI to produce feedback while demanding exhaustive disclosure from students, the asymmetry teaches its own moral lesson. Protestant suspicion of fallen institutions should apply to universities and professors as rigorously as to student behavior.

This doctrine also rules out complacency about “human oversight.” Humans can be biased, self-interested, fatigued, deferential, or captured by incentives. Keeping a person nominally in the loop does not guarantee wisdom. Governance must distribute authority, invite dissent, protect whistleblowing, record provenance, and make consequential decisions contestable. In settings where junior members may be reluctant to challenge senior authority, leaders have a special duty to create structured forms of disagreement rather than interpreting silence as consent.

4.4 Incarnation and embodied, relational knowing

Christian faith is not a message about escape from material existence. The Word became flesh and dwelt among us (John 1:14); Christ taught through meals, questions, journeys, touch, confrontation, suffering, and shared life. The resurrection affirms embodied life rather than replacing it with disembodied information. The cross, no less than the incarnation, keeps Christian anthropology from idealism: God addresses human finitude and sin not by abstracting persons into information but by entering embodied, historical particularity and bearing its full cost (Phil. 2:6–8). A theology of AI in education that draws only on incarnation as gentle divine accompaniment, without the cross's more costly claim that reconciliation required atonement, risks reducing Christology to a warrant for "presence" in the abstract. Section 4.8 returns to the cross directly, as the specific ground of the grace this paper invokes throughout. Incarnation should not be turned into a simplistic rule that every valuable educational interaction must occur face-to-face. It establishes a deeper claim: persons know and are formed as embodied beings whose understanding is situated in places, practices, histories, and relationships.

Much university knowledge is irreducibly tacit and enacted. A medical student learns how to attend to a patient; an engineer develops care around materials and failure; a field researcher learns what a dataset omits; a teacher notices confusion in a room; a laboratory apprentice acquires safe and responsible habits; a manager learns how a decision is received by those who bear its consequences. Textual explanation contributes to each activity but cannot replace participation in it. Generative AI is strongest precisely where universities are tempted to confuse an articulate representation of practice with the practice itself.

Confucian education offers a strong parallel emphasis on formation through relationships, exemplars, and ritualized practice. Teacher and student are not interchangeable nodes in an information network. Respect for teachers can sustain humility, attentiveness, gratitude, and continuity across generations. *Li* can also draw attention to the moral significance of ordinary educational forms: how one listens, cites, disagrees, asks for help, gives feedback, and acknowledges dependence (Lam 2026). These practices can preserve dimensions of education that disappear when learning is imagined as private content consumption.

Yet hierarchy must remain open to criticism. Respect can become unquestioning deference; harmony can silence legitimate conflict; seniority can protect error. A Christian teacher's authority is ministerial rather than absolute: it is exercised for the student's good and remains accountable to truth, neighbor-love, and the judgment of a wider community. Jesus' own pattern of authority joins teaching with service and refuses status as domination (Mark 10:42–45). AI may sometimes provide a useful alternative channel for a student who fears embarrassment or loss of face when asking a basic question. That is a real benefit, especially in large or strongly hierarchical classrooms. But the remedy should not be a university in which students can learn only by avoiding vulnerable contact with teachers and peers. Leaders should improve the human ecology as well as provide a digital refuge.

The contrast is not “cold machines” versus “warm humans.” Human teachers can be inattentive, humiliating, or unavailable; AI can appear patient, accessible, and responsive. The relevant difference is reciprocity and answerability. A teacher can know a student's history, be corrected, take responsibility for advice, and remain in a relationship whose obligations extend beyond the current prompt. A peer risks something in disagreement and collaboration. A language model simulates conversational responsiveness without sharing the vulnerability, consequences, or common life of the learner. It can support a relationship-rich pedagogy, but it cannot supply the whole good of educational relationship.

University leaders should therefore identify **presence-critical practices** rather than defend presence abstractly. Clinical encounters, laboratories, fieldwork, oral defence, studio critique, pastoral care, difficult ethical deliberation, and some forms of apprenticeship carry goods that arise through embodied participation and mutual accountability. AI may prepare students for these practices or help them reflect afterwards. It should not become evidence that the practices themselves are obsolete.

4.5 Truth, wisdom, and sincerity under conditions of fluent simulation

Christian education is ordered toward truth because creation is God's reality, not material for constructing whatever answer best serves the self. Yet human knowledge is partial and socially mediated. A Christian commitment to truth should therefore produce both confidence that truth matters and humility about one's own grasp of it. The learner must distinguish testimony from proof, rhetoric from reason, consensus from reality, and personal conviction from warranted judgment.

Generative AI sharpens this task because it can produce the linguistic marks of understanding without being a responsible knower. Its output may be correct, useful, and illuminating. But the system does not stand behind a claim as a scholar does: it does not love truth, confess error, risk reputation, or owe an account to the people affected. The learner remains responsible for the decision to trust, revise, reject, or communicate the output. Fluency is therefore neither evidence of understanding nor a transfer of epistemic responsibility.

Confucian learning provides an especially fruitful dialogue partner here. Analects 2:15 joins learning with reflection; *cheng* connects inward sincerity with outward expression; *zhi* concerns wisdom or discerning judgment. Lam argues that LLM-assisted work becomes morally problematic when polished expression is detached from the student's actual thought and when external information is not transformed through reflection (Lam 2026). This account deepens academic integrity beyond rule compliance. Undisclosed generation is wrong not only because it violates a regulation but because it fractures the relation between the presented work and the learner who asks others to receive it as his or her own.

Protestant theology both affirms and qualifies this account of sincerity. Truthfulness requires congruence between claim and authorship, but sincerity alone does not make a claim true. A student can be authentically mistaken, and a community can sincerely preserve error. The Reformation's insistence that institutions and traditions remain corrigible under the Word of God provides a resource against both algorithmic authority and human conformity. Respectful challenge is not the enemy of harmony when the apparent harmony is sustained by silence or falsehood.

This account of truthfulness remains incomplete without attention to the Spirit whom Scripture identifies as the one who "will guide you into all truth" (John 16:13) and who searches even "the deep things of God," granting believers a mind capable of discerning what has been freely given by God rather than merely what appears plausible (1 Cor. 2:10–14). Protestant theology has sometimes reduced this promise to private illumination in personal Bible reading. Its scope is wider: Paul lists the "discerning of spirits" among the gifts given for the building up of the whole community (1 Cor. 12:10), a gift exercised corporately, through testing what is said against Scripture and against the discernment of others, not through solitary certainty. Amos Yong's pneumatological account of discernment is instructive here: authentic discernment is neither naive trust in every plausible-sounding claim nor reflexive suspicion of every novel voice, but a disciplined, communal practice of testing (Yong 2000). Similarly, Simon Chan's account of spiritual theology likewise insists that discernment is formed through sustained participation in a worshipping community's practices, not derived from private intuition alone (Chan 1998).

This has direct bearing on the epistemic problem posed by fluent AI output. A generated answer is, in a loose but instructive analogy, a "spirit" in the sense Scripture uses that term for any voice claiming authority to speak truth: it must be tested, not simply received (1 Thess. 5:21; 1 John 4:1). A Christian university should therefore treat discernment of AI-generated claims not merely as a cognitive skill to be taught procedurally (Section 5.1) but as a practice formed by, and modeled on, the community's larger discipline of testing every truth-claim—including its own most confident ones—against Scripture, against the wisdom of a gathered community, and in humility before the Spirit who alone finally guides into truth. A pedagogy of AI verification that trains only the technical habits of citation-checking, while neglecting the formation of a community capable of prayerful, corporate, Spirit-attentive discernment, has quietly secularized a practice the New Testament treats as inescapably communal and spiritually formed.

The Asian linguistic context introduces a further epistemic responsibility. Students may work across Mandarin, English, regional languages, dialects, and disciplinary vocabularies. AI translation and drafting can widen participation and help students test how an idea travels between linguistic worlds. It can also homogenize expression, privilege dominant languages and sources, and make culturally unfamiliar formulations appear universally authoritative. UNESCO has warned that insufficiently localized models can produce culturally irrelevant or biased educational content in the Asia-Pacific (UNESCO 2023). Christian professors should therefore ask not only whether an output is factually accurate, but which archives, languages, histories, and communities are rendered visible or invisible by the system.

This concern is particularly important for Christian scholarship in Chinese contexts. A general-purpose model may generate an apparently competent answer from globally dominant English-language theological material while flattening the diversity of Chinese Christian experience, confusing Confucian concepts, or importing Western culture-war categories. Conversely, a locally optimized system may reproduce the omissions or preferred narratives of its own training environment. No model is culturally unlocated. Students should learn to compare outputs across languages, inspect sources, retrieve primary texts, consult locally situated scholarship, and explain where translation involves conceptual judgment rather than word substitution.

Wisdom finally exceeds verification. A model may correctly summarize the available options yet cannot determine what should be loved, which burden should be accepted, or how a general rule should be enacted faithfully in a singular relationship. Wisdom integrates knowledge, character, attention, consequences, and rightly ordered ends. It is cultivated through practice and correction, sought in prayer and community, and never converted into a confidence score. Educational AI is most defensible when it supplies material for this judgment rather than impersonating its completion.

4.6 Vocation and stewardship: from achievement to answerable service

The Protestant doctrine of vocation widens the meaning of calling beyond clergy and exceptional religious work. Teaching, engineering, medicine, computing, administration, parenting, citizenship, and ordinary forms of service can become sites of faithfulness. Knowledge and skill are gifts to be cultivated for God and neighbor, not private possessions or merely competitive assets. The priesthood of all believers also denies that social prestige maps neatly onto spiritual worth: visible leadership, technical expertise, care work, and uncelebrated forms of maintenance can each serve the common good.

This account resonates with Confucian resistance to the isolated self. Learning carries obligations to family and society. In Chinese and other Asian settings, the university degree often represents a collective investment, and a student's career decision may be experienced as consequential for parents and extended family. A Christian theology should not dismiss these obligations as external interference in an individual's authentic choice. Scripture also treats persons as members of families and communities, commands honor for parents, and directs gifts toward service.

But no human relationship is ultimate. Family honor, institutional loyalty, national development, professional status, and economic security are genuine goods that can become idols. The Christian calling may require disappointing an expectation, exposing institutional failure, refusing an unsafe deployment, or choosing service that carries less prestige. Reciprocal filial responsibility can support agency and gratitude; authoritarian obligation can make love conditional on performance. The empirical distinction identified by Zhou, Guo, and Xu (2020) is not itself a theological conclusion, but it supports a more careful pastoral and institutional account than either “Asian family values” or “individual autonomy.”

AI makes the purpose of expertise a pressing vocational question. China and other Asian economies are rapidly aligning curricula with AI-enabled industrial transformation. Preparing students for this world is part of responsible stewardship. Yet universities should not equate employability with immediate tool proficiency. A graduate's vocation requires sufficient disciplinary understanding to notice error, resist pressure, explain a decision, and remain answerable to those affected. If AI performs an analysis, writes code, or produces a diagnosis, the human professional cannot transfer moral responsibility to the system. Delegated execution may change the distribution of work; it does not erase the calling to judgment.

Stewardship also asks who benefits. Automating routine documentation may release a clinician to attend to patients or a professor to mentor students. It may instead reduce staffing, intensify workloads, and concentrate gains in the institution or vendor. An AI system may extend high-quality tutoring to

rural or lower-income students, yet paid tiers, infrastructure gaps, and uneven prior knowledge may distribute benefits upward. In the Asia-Pacific, differences in technical readiness are substantial (UNESCO 2023). Leaders must therefore examine the whole labor and educational system, not only the experience of the immediate user.

The vocational test for any AI-mediated assignment is consequently twofold: **What capability is the student learning to exercise, and for whose good will that capability be exercised?** A student should be able to state what was delegated, what judgment remained human, who could be harmed by error, and why the final action is defensible. This turns AI literacy from consumer competence into accountable stewardship.

4.7 Community, harmony, and the love that permits correction

Christian personhood is irreducibly relational. Human beings depend on testimony, care, correction, and gifts they did not create. The church's image as a body with many members rejects both isolated self-sufficiency and uniformity (1 Cor. 12:12–27). Love of neighbor is not generalized friendliness. It entails attention to a concrete other, willingness to bear obligations, truthful speech, forgiveness, and concern for the vulnerable.

This relational anthropology has substantial affinity with Chinese traditions that understand the person through networks of family and social relationship. It also challenges a common trajectory of educational technology: increasingly private interaction between an individual learner and a personalized system. AI assistance can make support available at any hour and may be especially valuable to a student who is linguistically hesitant, geographically remote, disabled, or reluctant to expose uncertainty publicly. Such accessibility can enable fuller participation in human community.

But personalization is not identical with belonging. A system can remember preferences and produce empathic language without becoming a member who owes fidelity, receives correction, or shares consequences. The more agreeable and frictionless the interaction becomes, the more ordinary human relationships—where others misunderstand, disagree, make demands, and cannot be reset—may appear defective. The sycophancy evidence reviewed in Section 3 gives this theological concern an empirical mechanism: users can prefer affirming responses even when those responses reinforce their own position and reduce intentions toward relational repair.

The concept of harmony requires particular care. At its best, harmony is not sameness but the fitting coordination of distinct voices and responsibilities. It can sustain patience, courtesy, and a commitment to the relationship through disagreement. At its worst, “harmony” becomes a demand that junior members suppress questions, students protect the teacher's face, or institutions conceal conflict. Christian love refuses both aggressive individualism and peaceful-looking falsehood. The New Testament pattern of speaking truth in love (Eph. 4:15) joins correction to belonging rather than making one the alternative to the other.

This has direct implications for pedagogy in hierarchical classrooms. Professors should create structured occasions in which students must question an AI output, a peer, and the professor; assign roles that

authorize junior participants to identify assumptions; reward revision after criticism; and distinguish respect from agreement. Anonymous or AI-mediated rehearsal may help students prepare a difficult question, but it should lead back into accountable human dialogue. University leaders should likewise include students, junior faculty, professional staff, disability advocates, and less technologically privileged groups in AI governance instead of consulting only senior leaders and technical experts.

Community is also the setting in which cultural and theological discernment becomes possible. An individual professor may not notice that a model marginalizes a language, assumes a Western family structure, or treats a locally contested value as universal. A diverse community can expose these limits. For Christian institutions, this community must include the lived faith of churches and ordinary believers, not only academic theology. Asian Christian theologians have repeatedly cautioned against contextual theology produced solely by elites and detached from grassroots Christian life. AI governance that is technically sophisticated but socially insulated risks reproducing the same error.

4.8 Grace and hope: freedom from meritocratic and technological salvation

The argument culminates in grace and hope because Christian anthropology is not merely a description of human capacities and duties. Human beings are reconciled to God by grace, not by intellectual achievement, moral self-cultivation, institutional excellence, or technological progress. Their future rests in God's faithfulness and the resurrection of Christ, not in perpetual optimization. This is the decisive Protestant interruption of both meritocratic anxiety and technological salvation.

This is where the Protestant tradition offers its most distinctive, and for the anxieties described throughout this paper its most needed, contribution. The Reformation’s central recovery was not a general sentiment that persons matter apart from performance; it was the specific claim that a sinner is reckoned righteous before God by faith in Christ’s finished work, apart from the works of the law (Rom. 3:21–26; Gal. 2:16; Eph. 2:8–9). Righteousness is received, not achieved—an alien righteousness, in Luther’s phrase, credited to the believer rather than accumulated by the believer’s own record. This is a sharper claim than “human worth does not depend on performance” treated as a general humanistic conviction; it is the specific gospel claim that the ultimate verdict on a person’s worth has already been rendered, in Christ, prior to and independent of any grade, credential, publication record, or AI-assisted achievement (McGrath 2005). A Christian university’s confidence that assessment need not function as a verdict on a student’s ultimate worth is therefore not a generically “Christian” humanistic sentiment borrowed for institutional comfort; it rests on a specific soteriological claim that this paper’s other theological resources—creation, the image of God, vocation—support but do not, by themselves, supply.

This matters concretely for the meritocratic anxiety documented in Chinese and other Asian educational contexts (Sections 4.2 and 4.6). Confucian self-cultivation, taken on its own terms, tends to locate moral worth in the achieved outcome of disciplined practice; grace, in the Protestant sense argued here, locates it in a gift received prior to any practice. The two need not be opposed in every particular—cultivation remains a faithful response to grace already given, as Section 4.1 argues—but the direction of dependence differs, and a Christian university should be candid about that difference rather than

blending the two accounts into an ambiguous appeal to “unearned dignity” that neither tradition would fully recognize as its own. Justification by faith alone gives Christian professors and students a specific, not merely a general, resource for facing an assignment, a viva, or a career setback without treating the outcome as a referendum on their worth before God.

Grace does not make effort irrelevant. It liberates effort from the burden of creating worth. A student may study diligently to understand, serve, and honor legitimate obligations without treating each result as a verdict on his or her value. A professor may pursue excellence without making citation counts, rankings, or technical relevance into an identity. A university may innovate without believing that institutional survival justifies every adoption. Failure can then become an occasion for truth, correction, and growth rather than a shame that must be concealed behind generated competence.

This claim has particular force where education is bound to family sacrifice and intense competition. Many Chinese and Asian students do not experience achievement as an individual possession; success may repay parents, secure family mobility, and preserve face. Those motives can express gratitude and love. They can also make an imperfect grade, uncertain career, or request for help feel like moral failure. AI offers an alluring means to prevent visible weakness. A grace-shaped university should not respond by condemning family obligation or lowering standards. It should cultivate communities in which students are valued before they perform, assistance can be sought without humiliation, honest failure is less threatening than concealed success, and excellence is directed toward service rather than self-justification.

Hope similarly judges rival eschatologies. Some AI narratives promise abundance without toil, mastery over biological limits, perfect personalization, or the solution of social conflict through sufficient prediction and optimization. National and institutional narratives may be less extravagant but structurally similar: technological leadership will secure the future, and universities must transform immediately or become irrelevant. Christian hope does not require pessimism about AI. It permits energetic, creative work because the final redemption of creation is God's work, not the university's engineering project (Langford 2022a). Leaders are therefore free to welcome a useful system, refuse a harmful one, revise a failed strategy, and admit uncertainty without imagining that history depends on uninterrupted technological advance.

Grace and hope also set a limit on formation language. University practices matter profoundly, but no curriculum, honor code, AI tutor, or Christian mission statement can manufacture saints. Institutions should design for truthful authorship, patient learning, responsible service, and human community while recognizing the opacity and freedom of persons. This humility guards against both surveillance and paternalism. The aim is not to make inward virtue legible to an algorithm; it is to establish practices in which responsible action is possible, expected, supported, and open to correction.

4.9 A Chinese-contextualized synthesis

The dialogue can be summarized as neither assimilation nor opposition. Confucian educational traditions provide moral and relational insights that can correct thin, individualistic approaches to AI ethics. Evangelical theology affirms many of these goods while locating dignity in creation, judging all

communities under the reality of sin, and grounding worth and hope in grace rather than cultivation or achievement.

Contextual theme	Constructive resonance	Characteristic danger under AI mediation	Evangelical contribution	Institutional implication
Learning and self-cultivation	Learning should transform the learner, not merely transfer information.	Generated fluency and credential completion can replace reflection and sincere intellectual labor.	Formation is a response to grace and calling, not a means of earning dignity.	Assess process, reflection, explanation, and retained competence, not only polished outputs.
Filial and social obligation	Education can express gratitude and responsibility to family and society.	Achievement pressure may make AI-assisted concealment attractive and bind worth to family honor.	Honor and service matter, but no family, institution, or nation is ultimate.	Provide confidential human support; teach vocation as answerable service; avoid shame-based integrity regimes.
Respect for teachers and seniority	Apprenticeship, gratitude, and receptivity to correction are educational goods.	Confident systems and senior leaders may receive deference without adequate challenge.	Authority is accountable and exercised in service; all human institutions are fallible.	Build structured dissent, appeal routes, oral defence, and reciprocal faculty disclosure of AI use.
Harmony and relationality	Learning occurs through relationships and responsibilities rather than isolated autonomy.	Harmony may become conformity; personalized AI may displace difficult human disagreement.	Truth and love belong together; community is differentiated rather than uniform.	Preserve peer critique and teacher presence; use AI to rehearse participation, not replace accountable dialogue.
Discipline, effort, and endurance	Repeated practice and patient reflection form competence and character.	AI can bypass formative struggle, while institutions use efficiency to intensify workload.	Creatureliness and Sabbath affirm limits, rest, and dependence without romanticizing needless burden.	Identify stage-specific AI-free practice; ensure efficiency gains return time to learning, rest, or relationship.
Educational and technological modernisation	AI literacy can widen capability and support social development.	Universities and students may be instrumentalized for rankings, labor-market throughput, or technological competition.	Technology is a form of stewardship, not salvation; the neighbor and common good judge innovation.	Evaluate distribution of benefit and risk, cultural localization, labor effects, and mission alignment.
Linguistic and cultural plurality	AI can translate, scaffold participation, and connect knowledge traditions.	Dominant languages and training data can homogenize thought or erase local and minority knowledge.	Pentecost imagines unity without erasing linguistic difference; truth requires situated, corrigible communities.	Test tools across local languages; teach cross-language source verification; include local scholars and communities in evaluation.

This synthesis yields a more demanding anthropology than either technological humanism or cultural conservatism. The student is neither an autonomous consumer nor merely a bearer of family and national expectations. The professor is neither an information provider nor an unquestionable moral exemplar. The university is neither a credential factory nor a redemptive community. Each is situated

within relationships and institutions, accountable for the use of gifts, vulnerable to distorted loves, and dependent on grace.

Formation-sensitive augmentation follows from this theological sequence. AI should be used in ways that honor unearned dignity, respect creaturely development, expose rather than exploit disordered incentives, preserve embodied and accountable relationships, strengthen truthful and sincere judgment, equip answerable service, support communities in which correction is possible, and resist every promise that technical performance can save the person or institution. Section 5 now applies this anthropology to six domains of formative impact.

5. Six domains of formative impact

The language of formation can become vague unless it identifies what is being formed, through which practices, and with what observable consequences. This section therefore distinguishes six domains in which habitual AI use may shape students and universities: epistemic, cognitive, moral, relational, vocational, and institutional. The domains are analytical rather than separable. A student who repeatedly submits claims whose provenance she cannot trace faces an epistemic problem, but also a cognitive loss of understanding and a moral problem of authorship. A university that automates student advising may gain institutional efficiency while altering relationships of care and the professional vocation of advisers. The point of the framework is not to assign every use to a single category, but to make such interactions visible.

The analysis also requires evidentiary restraint. Present research can establish some near-term effects on performance, learning, engagement, and behaviour; it cannot yet demonstrate the full formation of a person over years. The claims below therefore combine three registers that should not be confused: empirical findings about observed effects, interpretations of the practices those effects may normalize, and theological judgments about the goods at stake. “Risk” denotes a plausible formative direction requiring investigation, not an assertion that every use produces the same result.

5.1 Epistemic formation: from plausible language to answerable knowledge

The epistemic domain concerns how students come to know, what they count as evidence, how they respond to uncertainty, and whether they remain open to correction. Generative AI changes the environment of inquiry because it produces finished linguistic form before the learner has necessarily encountered the underlying reasons. Its fluency can compress distinctions among a documented claim, a reasonable inference, a contested interpretation, and an invented statement. The danger is deeper than occasional hallucination: students may learn to treat plausibility as a sufficient warrant for belief.

This danger is intensified when the system obscures provenance. Conventional scholarship allows a reader to move, however imperfectly, from a claim to a source, from a source to its method and context, and from these to a judgment about reliability. A generated synthesis may combine material without revealing which source supports which proposition. Even retrieval-enabled systems can cite

sources that do not entail their claims. If students learn to "check the citations" only by confirming that a linked document exists, they have not yet learned evidential verification.

The alternative is not to exclude AI from inquiry. AI can expose students to competing explanations, generate counterarguments, translate unfamiliar terminology, help them formulate questions, and provide material for error detection. These benefits become epistemically formative when the student must reconstruct rather than merely receive the warrant for a claim. A productive task might require students to classify each sentence in an AI response as supported, inferred, disputed, or unsupported; identify what evidence would change the conclusion; compare the response with primary sources; or write the strongest case against the model's recommendation. In such uses, the machine's fluency becomes an object of disciplined judgment rather than a substitute for it.

Conversational sycophancy makes this issue relationally and morally salient. Cheng et al. (2026) found that sycophantic responses can increase confidence in one's own position and reduce prosocial intentions even while users prefer the interaction. For Christian education, correction is not an inconvenience imposed on learning; it is one of the means by which pride is confronted and wisdom grows. A system that reflexively validates the user can make epistemic humility less habitual. Professors should therefore teach students to elicit disagreement, identify disconfirming evidence, and distinguish emotional affirmation from intellectual warrant.

Asian and Chinese settings add two complications. First, AI may reduce the linguistic disadvantage experienced by students working in English or another non-native academic language. Translation, vocabulary support, and alternative explanations can enable genuine participation. At the same time, English-dominant training data and easily available global sources may displace Chinese-language scholarship, regional cases, minority languages, and local forms of knowledge. Translation can also conceal conceptual differences by mapping culturally situated terms onto superficially similar English ones. Second, where respect for teachers is expressed through reluctance to challenge authority, a confident chatbot may become another voice to which students defer. Private use can sometimes lower fear of losing face and make it easier to ask elementary questions; it must nevertheless lead back into communities where claims can be questioned publicly and responsibly.

Evidence of healthy epistemic formation would include students who can trace claims to appropriate sources, state uncertainty, distinguish evidence from inference, correct an AI answer, and explain why they accept a conclusion. Warning signs include polished work with weak source comprehension, fabricated or decorative citations, uniform arguments across students, inability to defend claims orally, and growing confidence without corresponding disciplinary competence.

5.2 Cognitive formation: extending thought without bypassing its development

The cognitive domain concerns attention, memory, conceptual understanding, problem solving, metacognition, creativity, and transfer. All intellectual work involves selective offloading. Writing stabilizes memory; diagrams externalize relations; calculators remove routine arithmetic; search engines extend access to information. The relevant question is not whether cognition is distributed across

people and tools, but which internal capabilities a learner must first acquire and continue to exercise in order to use those tools intelligently.

The evidence reviewed in Section 3 establishes the possibility of a performance–learning divergence. Structured AI tutoring can improve immediate learning, especially when it asks questions, provides calibrated hints, and adapts explanations. Unrestricted answer generation can improve work completed with the tool while producing no advantage—or a disadvantage—when students must later work independently (Bastani et al. 2025; OECD 2026). These findings support a developmental principle: assistance should normally follow the structure of learning, not merely the structure of task completion.

Novices and experts therefore should not be governed identically. An expert programmer may use a coding assistant to recall an interface while retaining a robust model of data structures, control flow, testing, and security. A novice who generates an entire program may acquire only the ability to recognize a plausible-looking solution. An experienced scholar can use AI to explore alternative outlines because she can reject a weak conceptual structure; a first-year student may accept the outline before learning what makes an argument coherent. The apparent democratization of expert output can conceal unequal development of expert judgment.

Productive struggle is important here, but it should not be romanticized. Difficulty forms capability when it requires the learner to retrieve, discriminate, integrate, revise, or persist toward an intelligible goal. Confusing instructions, inaccessible materials, repetitive formatting, and unnecessary linguistic barriers need not be preserved in the name of character. Formation-sensitive teaching distinguishes **generative difficulty**, which exercises a target capability, from **extraneous difficulty**, which consumes effort without serving the learning objective. AI can appropriately reduce the latter while protecting the former.

For example, a multilingual engineering student may use AI to clarify the language of a design brief without delegating the mathematical modelling. A dyslexic student may use speech-to-text while retaining responsibility for the argument. A student learning optimization may use AI to explain a solver error after first specifying the decision variables and constraints, but should not submit a model whose feasibility assumptions cannot be defended. These distinctions prevent independent competence from becoming a proxy for a narrow, unaided mode of expression.

The Chinese tradition of disciplined practice offers a genuine resource: mastery often requires repetition, imitation of good models, correction, and gradual internalization. Yet practice loses its formative function if AI supplies the answer before the student attempts retrieval or construction. Conversely, intense examination cultures can produce large quantities of repetitive work with limited understanding. AI can help vary examples, diagnose misconceptions, and provide low-stakes practice, provided it does not simply accelerate an already over-optimized regime. The objective is neither maximum difficulty nor maximum ease, but well-ordered practice.

Healthy cognitive formation is visible when students can perform core operations after assistance is removed, transfer knowledge to unfamiliar cases, detect model errors, explain their strategy, and judge when they need help. Warning signs include dependence on prompts for basic initiation, declining recall

of foundational concepts, inability to adapt generated solutions, shallow confidence, and a widening gap between the sophistication of submitted work and the student's independent explanation.

5.3 Moral formation: truthfulness, effort, authorship, and responsibility

Every assessment regime teaches a moral lesson. It communicates what counts as honest work, whether reasons matter or only answers, whether rules are intelligible, and whether responsibility can be displaced onto a tool. AI policy therefore forms character not simply by prohibiting misconduct but by structuring repeated practices of disclosure, verification, attribution, correction, and answerability.

Academic integrity remains important, but a purely detection-centered response is inadequate. AI detectors are unreliable and can disadvantage non-native English writers (Liang et al. 2023). More fundamentally, surveillance can produce a compliance culture in which students ask only what can be detected. Ambiguous rules create a different injustice: students with greater cultural confidence, technical knowledge, or access to private advice learn how to exploit uncertainty, while more cautious students forgo legitimate assistance. Clear, task-specific expectations are therefore a condition of fairness as well as enforcement.

Authorship should be understood substantively. A student is not the author merely because her name appears on the submission or because she edited generated sentences. Authorship entails being able to endorse the claims, account for the process, identify material assistance, and answer criticism. This does not require that every word originate unaided. Scholarship has always involved editors, collaborators, software, and sources. It requires that assistance not conceal whose judgment governs the work.

The moral significance of effort also needs precision. Christianity does not sanctify inefficiency, and grace does not make achievement irrelevant. Some labour should be reduced because it is tedious, exclusionary, or poorly connected to the educational good. Yet the habitual desire to eliminate every demanding step can train avoidance. Students may gradually treat patience, revision, and uncertainty as system failures rather than ordinary features of learning. Formation-sensitive practice asks whether the effort being removed is accidental burden or the very practice through which attention, courage, honesty, and responsibility are learned.

In many Chinese and other Asian settings, high-stakes performance is entangled with family sacrifice, institutional reputation, and the preservation of face. These obligations can motivate admirable diligence, but they can also make disclosure of difficulty feel shameful and generated competence attractive. A punitive policy may drive use underground. Professors should create legitimate avenues for help, make disclosure ordinary rather than humiliating, and separate confession of uncertainty from loss of worth. The Protestant Christian claim that dignity and grace precede achievement can challenge both individual deception and systems that make failure unbearable.

AI also creates the possibility of **moral deskilling**. In professional cases, students may ask a system not only to calculate or summarize but to decide what is fair, safe, or humane. The danger is not that the model contains no moral language; it can generate abundant ethical vocabulary. The danger is that the student does not learn to perceive affected persons, weigh conflicting goods, or accept ownership of the

decision. AI may enrich deliberation by surfacing stakeholders and alternative arguments, but final moral judgment must remain answerable to persons and communities who can call the decision-maker to account.

Healthy moral formation appears in candid disclosure, accurate attribution, willingness to revise, proportionate reliance on assistance, and readiness to accept responsibility for consequences. Warning signs include concealed dependence, formulaic ethics sections appended to technical work, blaming the tool for errors, strategic compliance with unclear rules, and a culture in which output quality excuses misrepresentation.

5.4 Relational formation: presence, correction, and the limits of simulation

University learning is mediated by relationships: novice and expert, student and peer, researcher and community, professional and client, teacher and inherited tradition. These relationships do more than transfer information. They provide recognition, correction, example, trust, obligation, and opportunities to become responsible for others. AI can participate in communication, but it cannot occupy these roles in the same moral sense because it cannot be vulnerable to the consequences of advice, owe fidelity to a student, or share membership in a community of accountability.

The constructive possibilities are nevertheless substantial. An always-available tutor can answer questions that a student is embarrassed to ask, offer additional practice, support learners across time zones, and make materials more accessible. Students who are shy, disabled, isolated, or studying in a second language may gain a lower-risk entry point into learning. Faculty with heavy loads may use AI to produce preliminary feedback while preserving time for higher-value interaction.

The formative question is whether such uses become a bridge into human participation or an alternative to it. If students consult a chatbot rather than risk disagreement with peers, ask for affirmation rather than receive correction from a mentor, or simulate a stakeholder instead of meeting one, convenience can narrow the relational world of education. Simulated empathy may feel attentive while requiring nothing from the user in return. The student can terminate, regenerate, or redirect the conversation without learning patience with an independent other.

This issue is especially nuanced in relational and hierarchical cultures. Respect for teachers and concern for group harmony can sustain apprenticeship and communal responsibility. They can also inhibit questions, disagreement, and the disclosure of confusion. AI can create a private rehearsal space: a student may test a question, translate it, or consider a counterargument before speaking in class. A good design uses this confidence to increase participation. A poor design allows the student—and perhaps an overburdened institution—to avoid the human encounter indefinitely. Genuine harmony is not the absence of challenge; it is a relationship strong enough to sustain truthful correction.

Professors should therefore identify **relationally non-delegable moments**. These include conversations about vocation or serious failure; assessment of a student's readiness for consequential professional practice; feedback in which identity and confidence are at stake; resolution of conflict; pastoral care;

and engagement with people affected by a design or policy. AI may help prepare for these encounters or document agreed actions. It should not become the responsible party within them.

Healthy relational formation is visible when AI-supported students participate more fully with teachers and peers, seek correction, collaborate without hiding contributions, and encounter real stakeholders. Warning signs include declining attendance or consultation, substitution of simulated users for field engagement, avoidance of disagreement, dependence on affirming dialogue, and institutional reduction of human contact for those students who most need it.

5.5 Vocational formation: capability ordered toward answerable service

Universities must prepare graduates for workplaces in which AI will be normal. Refusing to teach AI use may leave students unprepared; equating employability with tool fluency may leave them unable to recognize when the tool fails. Vocational formation therefore requires a double competence: graduates should be able to work capably with AI and retain the disciplinary understanding, practical judgment, and moral responsibility needed to govern its use.

The content of retained competence differs by profession. A software engineer need not write every routine function from memory, but must understand architecture, testing, security, data provenance, and failure modes well enough to take responsibility for deployed code. An engineer may use AI to generate design alternatives, yet must be able to justify assumptions, safety margins, environmental effects, and trade-offs to regulators and affected communities. A clinician may use AI-supported decision tools, but must integrate patient history, uncertainty, consent, and values rather than treating a recommendation as a verdict. A business graduate may automate market analysis while remaining accountable for misleading data, discriminatory segmentation, or consequences shifted to workers and consumers.

The dividing line is not “human does the important work, AI does the routine work,” because systems increasingly participate in complex analysis. It is **humanly retained answerability**. A person or accountable group must possess sufficient epistemic control to understand the basis and limits of a recommendation, authority to reject it, and responsibility to explain and bear the consequences. Human oversight that consists only of clicking approval after an opaque process is ceremonial, not substantive.

Vocation also concerns ends. In a market-oriented university, students may understand AI primarily as a means to individual competitiveness. In national development discourse, including contemporary Chinese initiatives to embed AI literacy across education, AI competence may be framed as service to industrial upgrading and collective prosperity. These are legitimate goods, but neither market success nor national capability is self-validating. Christian vocation asks whom the work serves, whose burdens it increases, whether it contributes to justice and peace, and whether professional power is exercised as stewardship. Confucian concern for socially responsible learning can be a partner in this critique, while Christian theology also resists making family, corporation, or nation the final object of loyalty.

Healthy vocational formation is visible when students can use contemporary tools, perform essential checks, communicate uncertainty, refuse unsafe or unjust applications, and explain how their work

serves affected people. Warning signs include curricula organized around transient platform features, graduates who can prompt but not diagnose, ethics detached from technical decisions, and employers transferring responsibility to systems while retaining the benefits of automation.

5.6 Institutional formation: the university shaped by what it automates

AI does not form only individual users. Procurement, workload models, performance indicators, platform defaults, and assessment policies form the university as an institution. They determine which practices are funded, which relationships are treated as costly, whose knowledge is visible, and who has authority to define educational quality. An institution may profess relational education while deploying automation chiefly to reduce human contact. It may praise academic freedom while requiring faculty to use a single vendor. It may invoke inclusion while placing privacy and error risks disproportionately on students with the least power.

Institutional benefits are real. Shared services can provide accessible materials, translation, tutoring, research assistance, and relief from routine administration. Smaller institutions may gain capabilities previously available only to wealthy universities. Faculty may recover time for mentoring and scholarship. Yet these benefits are not automatic. Savings may be captured without reinvestment in educational relationships; centralized tools may standardize pedagogy; data collection may normalize surveillance; and dependence on a vendor may weaken institutional control over curriculum, research data, and intellectual property.

AI strategy therefore entails questions of political economy. Who owns the platform and the data? Whose labour produced the training material? Which languages and disciplines receive reliable support? Can students refuse a tool without educational disadvantage? Are errors contestable through meaningful human review? Does a pilot measure retained learning and equity, or merely usage and satisfaction? What happens if pricing, model behaviour, data location, or terms of service change? These questions are not peripheral compliance matters. They shape the conditions under which teachers and students can exercise agency.

Asian universities operate across very different political, linguistic, religious, and regulatory environments, so no single “Asian model” follows. Some institutions can coordinate infrastructure and curricular reform rapidly; such capacity may support equitable access and systematic faculty development. The same centralization can reduce local pedagogical discretion or make questioning institutional strategy difficult. China’s 2026 AI + Education Action Plan illustrates the scale and seriousness with which AI literacy and interdisciplinary education may be pursued. Formation-sensitive leadership should welcome investments in capability while asking whether literacy includes criticism of AI’s limits, alternative conceptions of human flourishing, and space for disciplines and communities not readily measured by national or market productivity.

For Christian institutions, mission drift can occur in two directions. A university may adopt commercial definitions of innovation and then add theological language after the fact. Or it may invoke Christian identity to resist change without investigating how AI could widen access and service. Faithful governance requires mission to shape the criteria of adoption from the beginning: truthfulness, dignity,

justice, stewardship, community, care for the vulnerable, and hope that is not dependent on technological success. Christian leaders in secular universities can advocate comparable practices through public concepts such as educational validity, human agency, professional accountability, privacy, accessibility, and the common good.

Healthy institutional formation is visible when governance remains contestable, faculty and students participate meaningfully, benefits are distributed fairly, human relationships are protected, and pilots can be revised or stopped. Warning signs include adoption targets without learning goals, compulsory dependence on one platform, reduction of faculty judgment to content delivery, disproportionate surveillance of vulnerable groups, and efficiency savings achieved by removing the relationships the institution claims to value.

5.7 Reading the six domains together

The six domains provide a diagnostic map rather than six independent scores. A proposed use should be examined for direct effects and for second-order interactions. For example, an AI feedback system may strengthen cognition through rapid practice but weaken epistemic formation if its explanations lack sources; it may improve access while weakening relationships if it replaces rather than supplements meetings with teachers. The proper response may be redesign rather than simple acceptance or rejection.

Domain	Central risk	Constructive possibility	Observable question
Epistemic	Plausibility replacing truth; opaque provenance; habitual deference	Comparison, explanation, critique, multilingual access	Can students trace, test, qualify, and defend the claims they use?
Cognitive	Offloading before competence; weakened recall, transfer, and metacognition	Adaptive tutoring, feedback, examples, deliberate practice	Can students retain and transfer the target capability when assistance changes or is removed?
Moral	Concealed dependence, avoidance of formative effort, diffusion of responsibility	Disclosure, attribution, ethical deliberation, accountable use	Do practices reward truthfulness and ownership rather than merely compliant outputs?
Relational	Displacement of teachers, peers, and affected communities; simulated affirmation	Accessible support and preparation for fuller participation	Does AI lead learners back into human correction, reciprocity, and care?
Vocational	Deskilling; performance detached from professional answerability	New capability, reduced routine work, wider participation	Can graduates understand, challenge, explain, and take responsibility for AI-supported work?
Institutional	Surveillance, inequality, vendor dependence, mission drift, deprofessionalization	Shared infrastructure, accessibility, administrative relief, research support	Who gains, who bears risk, and who retains authority over educational purposes?

The map changes the burden of proof. It is insufficient to show that an AI tool saves time or improves an immediate output. Nor is it sufficient to name a possible formative risk and demand prohibition. Advocates should specify the educational good, mechanism, safeguards, and evidence of benefit; critics

should identify the capability or relationship at risk, the plausible mechanism of loss, and whether redesign could mitigate it. Section 6 turns this reciprocal burden into a decision framework.

6. Formation-sensitive augmentation

Formation-sensitive augmentation is a principle of alignment: the scope and form of AI assistance should be aligned with the human capability, relationship, and responsibility that an educational practice exists to develop or demonstrate. It rejects three inadequate rules. “Use AI whenever it improves the output” confuses performance with education. “Never delegate cognitive work” ignores the tool-mediated character of expertise and the legitimate value of reducing extraneous burden. “Leave every decision to individual preference” ignores assessment validity, power, equity, and institutional responsibility.

The framework instead distinguishes **execution** from **answerability**. Execution includes producing text, code, calculations, summaries, alternatives, or predictions. Answerability includes understanding the basis of the work, evaluating its adequacy, disclosing relevant assistance, explaining the judgment made, responding to criticism, and bearing responsibility for foreseeable consequences. Execution may often be shared with or substantially delegated to AI. Answerability cannot be transferred to a system that cannot enter reciprocal moral relationships or be held responsible.

6.1 Begin with purpose, not the tool

An AI rule should follow from the purpose of the activity. Professors can begin by completing a simple sentence: “Students do this task so that they learn or demonstrate the capacity to ___.” If the answer is “produce a polished report,” the learning outcome remains underspecified. The relevant capacity may be framing a research question, evaluating evidence, constructing an argument, communicating to a client, or integrating these abilities. Only after the capacity is named can the instructor decide which forms of assistance support it.

The same logic applies institutionally. “Increase AI adoption” is an implementation target, not an educational purpose. A university may seek to widen access to tutoring, prepare graduates for AI-mediated work, reduce administrative delays, or improve research support. Each purpose implies different users, safeguards, measures, and possible alternatives. Naming the purpose also makes non-adoption a legitimate outcome when the proposed tool does not serve it.

6.2 Four determinants of appropriate delegation

The appropriateness of AI assistance depends on at least four interacting determinants.

19. **The target capability.** Is the activity meant to form factual recall, conceptual understanding, disciplinary analysis, communication, ethical judgment, collaboration, or professional performance? Assistance that is harmless for formatting may be destructive if the target is learning to structure an argument.

20. **The learner's developmental stage.** Does the learner possess sufficient prior knowledge to evaluate the output? Novices normally require more protected practice in foundational operations; advanced learners may appropriately automate some of them to address more complex problems.
21. **The system and interaction design.** Does the AI give answers immediately, or does it diagnose, question, hint, request justification, cite evidence, and expose uncertainty? General-purpose chatbots and pedagogically constrained tutors should not be treated as equivalent merely because both use a language model.
22. **The accountability context.** Is the task low-stakes practice, collaborative exploration, summative certification, research, or a decision affecting real people? The greater the stakes and irreversibility, the stronger the requirements for provenance, independent checks, human review, and retained answerability.

These determinants explain why universal lists of permitted tools age badly. A tool's capabilities change; the educational purpose remains the more stable point of reference. They also prevent "AI literacy" from becoming permission to use AI everywhere. Learning to use AI includes learning when not to use it, when to delay its use, and when its outputs must be subordinated to other forms of knowledge.

6.3 The ten-question governing test

For a proposed use—whether at the level of a task, course, programme, or institutional pilot—the following questions should be answered in writing.

23. **Purpose:** What educational good or institutional purpose is this activity meant to serve? A defensible answer names more than speed, adoption, or output quality.
24. **Capacity:** Which human capability, habit, relationship, or form of judgment is intended to develop or be demonstrated? Which parts are foundational and which are legitimately tool-extended?
25. **Delegation:** What precisely will the system generate, recommend, decide, or record? What must the student or professional initiate, verify, revise, and own?
26. **Developmental stage:** What prior knowledge does responsible use presuppose, and how will that knowledge be established? Would the same assistance be appropriate for a novice and an advanced learner?
27. **Epistemic control:** Can users inspect relevant sources, assumptions, reasoning, uncertainty, and limitations? What independent checks are required when the model's internal process is not inspectable?
28. **Responsibility:** Who signs the work, explains the judgment, responds to affected persons, and bears consequences? Does human review have the competence, time, and authority to reject the output?
29. **Relational effect:** Which interactions with teachers, peers, clients, patients, or communities will be increased, changed, or displaced? Is AI a bridge into participation or a substitute for presence?
30. **Distribution:** Who gains access and capability, and who bears risks of error, surveillance, dependence, language exclusion, or cost? Is a genuinely equivalent non-AI route available where appropriate?
31. **Evidence:** What observation would show genuine learning or mission advancement? Will evaluation examine retention, transfer, judgment, and equity rather than only completion time and satisfaction?
32. **Reversibility:** What are the review date, stopping conditions, data-exit provisions, and recovery plan if the practice produces deskilling, inequity, or mission drift?

The test is not an arithmetic score. A severe failure of responsibility, privacy, or assessment validity cannot be cancelled by several minor benefits. Its purpose is deliberative: to expose assumptions, identify redesign options, and create an accountable record of why a practice was adopted.

6.4 Four governance postures

After applying the test, a university or professor can adopt one of four explicit postures.

Posture	Appropriate when	Typical conditions
Require	AI competence is itself a learning outcome or the tool enables an essential pedagogical design	Equal access; instruction in limitations; disclosure; assessment of both tool use and retained judgment
Permit	Assistance does not displace the target capability and risks are low or well controlled	Clear scope; verification; attribution where material; accessible alternative if needed
Restrict or stage	Assistance could bypass novice formation, distort assessment, or create material risks	Tool-free or limited first attempt; approved functions; process record; independent checkpoint; oral defence
Prohibit for the activity	The task exists to establish unaided competence, involves unacceptable data or safety risk, or requires a relationship that cannot be delegated	Narrow, explained prohibition; feasible enforcement; appropriate support; review as tasks and tools change

A mature policy will use all four. Requiring AI in every course is as insensitive to purpose as banning it in every course. Restrictions should be task-specific and educationally explained. Permission should not mean absence of standards. Requirements should include equitable access and should never compel students to surrender sensitive personal or community data to a commercial system without adequate protection.

6.5 A provisional delegation matrix

Type of activity	Presumption	Formation-sensitive condition
Routine formatting, transcription, accessibility support, low-level administration	Generally augment	Confirm that the activity is not itself the learning target; protect privacy and provide equivalent access
Examples, hints, formative feedback, practice questions, alternative explanations	Augment with pedagogical constraints	Prefer questioning and graduated hints; require learner response; check accuracy and retained learning
Brainstorming, translation, coding assistance, preliminary synthesis	Context-dependent	Match use to learner stage; disclose material assistance; verify outputs; preserve local-language and disciplinary sources
Core analysis in early-stage learning	Restrict or stage carefully	Require an initial attempt or foundational demonstration before assistance; assess reasoning, not only the product

Type of activity	Presumption	Formation-sensitive condition
Summative demonstration of independent competence	Restrict unless AI competence is itself assessed	Specify the allowed boundary; include secure or dialogic checkpoints sufficient to support valid certification
Complex professional analysis	Permit only with strong controls	Require domain-specific validation, uncertainty communication, auditability, and authority to reject the system
Moral judgment, final authorship, and accountability for consequential decisions	Humanly retain	AI may inform deliberation, but identified people must understand, decide, explain, and answer to those affected

"Humanly retained" does not mean tool-free. It means that a person or accountable community remains the locus of judgment and answerability. A clinician can be answerable for an AI-supported diagnosis; a student cannot be answerable for an essay whose sources and reasoning she does not understand.

6.6 Testing the framework in four university practices

Academic writing. Suppose a first-year course uses essays to form question framing, close reading, argument, and revision. Allowing a chatbot to generate the initial thesis and structure risks bypassing the activity's central cognitive work. A formation-sensitive design might require students to annotate sources and submit an initial argument map before using AI. They could then ask the system for counterarguments, compare two revisions, document which suggestions they rejected, and defend the final position orally. In an advanced professional course where writing is primarily communication and the disciplinary analysis has already been demonstrated, more extensive language assistance may be appropriate. Multilingual students may use translation or language correction without surrendering authorship, but must check whether culturally and technically significant concepts have been altered.

Programming. In an introductory course, students need mental models of state, control flow, abstraction, data structures, testing, and debugging. Whole-solution generation can produce code that runs while concealing the learner's lack of understanding. Staged use might require prediction of program behaviour, manual construction of core functions, explanation of generated code, adversarial testing, and repair without the assistant. In an advanced software-engineering course, using a coding assistant may itself be required, because students must learn specification, review, security analysis, integration, and accountability for AI-produced components. The learning outcome shifts from unaided syntax production to governed engineering, but foundational competence is still sampled independently.

Engineering design. AI can generate alternatives, estimate parameters, and accelerate documentation. Yet design education should form judgment about assumptions, safety, sustainability, trade-offs, and affected communities. Students might use AI to expand the design space, but must maintain a decision log, validate calculations against authoritative methods, conduct failure-mode analysis, identify who bears residual risk, and present the design to human reviewers. A generic "ethical considerations"

paragraph is insufficient. Ethics belongs inside material choice, safety margins, accessibility, environmental impact, and deployment decisions.

Professional decision-making. In medicine, social work, law, business, or public policy, simulated cases can help students compare recommendations and expose hidden assumptions. They cannot fully substitute for accountable encounter with patients, clients, workers, or communities. A formation-sensitive case requires students to distinguish facts from model inference, state uncertainty, identify whose testimony is absent, and decide what must be discussed with the affected person. The final recommendation should be attributable to a named human decision-maker who can explain why the AI was followed or rejected.

Across all four practices, the framework asks the same question in discipline-specific form: does AI enable the learner to undertake more responsible intellectual work, or merely enable a more impressive artifact to appear?

6.7 A cycle of adoption, evidence, and correction

Formation-sensitive augmentation should operate as a cycle rather than a one-time approval:

33. specify the educational purpose and protected human capacities;
34. classify the proposed use as required, permitted, staged, or prohibited;
35. design access, disclosure, verification, relational, and accountability safeguards;
36. establish a baseline of independent capability where relevant;
37. pilot at a scale that permits close observation;
38. test assisted performance and later retention or transfer;
39. gather differentiated feedback, especially from students with less power or linguistic advantage;
40. revise, expand, restrict, or discontinue the use.

This approach is deliberately more demanding than announcing principles. It allows innovation, but makes educational benefit an empirical question and mission alignment an accountable judgment. It also embodies a Protestant posture of institutional humility: universities should not assume that good intentions or Christian identity immunize their technologies from distortion. Practices require continuing examination and reform.

7. Recommendations for professors

Professors are the most important interpreters of formation-sensitive augmentation because the educational meaning of an AI use depends on disciplinary purpose and student development. They should not be left to manage the transition without institutional support, nor should an institution replace their judgment with a universal tool policy. The following recommendations translate the framework into course practice.

7.1 Name the human capability behind each major task

Before deciding an AI rule, specify what students are meant to learn or demonstrate. Replace outcomes such as “write a report” or “produce working code” with capacities such as evaluating evidence, constructing a causal argument, modelling constraints, testing edge cases, communicating risk, or integrating stakeholder perspectives. Then identify which steps exercise those capacities.

This exercise often reveals that a single assignment serves several purposes that AI affects differently. A laboratory report may assess experimental design, statistical reasoning, scientific writing, and formatting. Language assistance might be permitted; fabrication of interpretations cannot be. Where the purposes cannot be cleanly separated, the instructor should stage the assignment rather than issue an all-or-nothing rule.

7.2 Design a progression from protected practice to governed use

Students need both foundational competence and the ability to work with contemporary tools. Courses should therefore sequence three kinds of experience where appropriate:

41. **without AI**, to acquire and demonstrate foundational concepts or operations;
42. **with pedagogically constrained AI**, to receive hints, questions, feedback, and additional practice; and
43. **with general-purpose or professional AI**, to learn verification, integration, limitation, and accountable use.

The sequence need not be identical in every discipline, but it prevents two failures: graduating students who have never learned to use tools normal in their profession, and graduating students whose apparent competence disappears when the tools fail or the context changes. The OECD’s 2026 recommendation to develop valued knowledge and skills without GenAI, with educational GenAI, and then with general-purpose GenAI provides convergent public-policy support for this staged approach.

7.3 Separate assistance-rich learning from valid checkpoints of competence

Not every activity needs to be secured against AI. Low-stakes practice can permit extensive experimentation if students receive timely correction. But every programme needs credible evidence of what students themselves understand and can do. Professors should identify a small number of **assurance checkpoints** aligned with the course’s essential outcomes.

These checkpoints can include oral explanation, live problem solving, supervised practical work, personalised application, critique of an unfamiliar output, laboratory performance, or a short tool-limited component combined with an AI-enabled project. The goal is not to recreate an entirely pre-AI examination regime. It is to triangulate evidence so that the credential remains truthful. TEQSA’s assessment-reform work similarly emphasizes learning assurance alongside responsible AI use; formation-sensitive augmentation adds the question of which habits and relationships the assessment process cultivates.

7.4 Assess process when process is educationally material

When learning occurs through framing, searching, drafting, debugging, testing, or revision, collect proportionate evidence of that process. Useful artifacts include a problem formulation, source trail, decision log, version history, test suite, error analysis, rejected alternatives, disclosure statement, and short oral defence. Students should explain not only what the AI contributed but why they accepted, modified, or rejected it.

Process evidence should be educational rather than forensic. Requiring thousands of words of prompts can create performative compliance and expose private data without showing understanding. Ask instead for decisions that reveal judgment. A concise ledger might record: task delegated; significant output used; verification performed; changes made; and remaining uncertainty. The burden should match the stakes of the task.

7.5 Teach AI literacy inside the discipline

Generic warnings about hallucination are insufficient. Students need to know how AI failure appears in their field. Historians should examine invented provenance and anachronism; lawyers, fabricated authority and jurisdictional error; programmers, insecure dependencies and tests that merely confirm generated assumptions; engineers, unjustified parameters and hidden safety trade-offs; clinicians, distribution shift and missing patient context; business students, biased segmentation and false precision.

UNESCO's student framework rightly combines a human-centred mindset, ethics, technical understanding, and system design, while its teacher framework integrates AI foundations with pedagogy and professional learning (Miao, Shiohira, and Lao 2024; Miao and Cukurova 2024). At university level these competencies should be translated into disciplinary epistemology: what counts as evidence, how uncertainty is expressed, which sources are authoritative, what must be independently checked, and who is answerable for error.

7.6 Use AI as an object of critique, not only a source of answers

Some of the strongest learning designs begin with an imperfect AI output. Students can compare responses produced from differently framed prompts, audit citations, diagnose a flawed proof, locate bias in a case analysis, improve a test suite, or identify the absent stakeholder in a recommendation. They can ask the system to defend opposing positions and then evaluate whether the disagreement is substantive or merely rhetorical.

Such tasks reverse the normal relation of deference: the student must exercise judgment over the output. But critique itself should not become a ritual in which students find a token mistake and then accept the rest. Rubrics should reward the quality of verification, the reconstruction of reasons, and the student's improved alternative.

7.7 Protect novice practices that function as modes of thought

Writing, coding, drawing, calculating, translating, and designing are not merely ways to display a prior thought; in many disciplines they are how thought develops. Professors should identify which forms of production are cognitively generative for novices and delay automation of those portions. This protection can be temporary and targeted.

The criterion is not nostalgia. Hand calculation may be essential until a student understands scale and structure, then unnecessary for advanced modelling. Writing an initial paragraph may be essential for discovering an argument, while later language editing may not be. Coding a data structure may matter once; repeatedly writing boilerplate may not. Professors should explain these distinctions so that restrictions are understood as part of a developmental pathway, not as arbitrary distrust.

7.8 Bring students into contact with real persons and consequences

Professional education should include tasks in which students encounter those affected by their analysis: users, patients, clients, workers, community partners, or practitioners. AI-generated personas and simulated interviews can prepare students for these encounters, expand scenarios, or reveal questions they have overlooked. They should not routinely replace actual testimony.

Require students to document whose interests were considered, which voices were directly heard, where conflicts arose, and how the final design or recommendation changed. In engineering and computing, ethical formation becomes concrete when students must defend data choices, failure handling, accessibility, environmental cost, and residual risk—not merely append an ethics section. In business and public policy, it requires attention to how efficiency gains and harms are distributed.

7.9 Make disclosure clear, ordinary, and non-humiliating

Every substantial assignment should state what AI use is required, permitted, restricted, or prohibited; why; and what disclosure is expected. Examples of acceptable and unacceptable use are more helpful than broad references to “responsible AI.” Rules should address translation, grammar correction, brainstorming, coding, image generation, synthesis, and citation assistance where relevant rather than assuming students know how each is classified.

Disclosure should be normalized as a practice of authorship, not treated as an admission of suspicion. This is especially important where preserving face or meeting family expectations can make acknowledgement of assistance feel shameful. Students should be able to disclose uncertainty and seek help without being treated as deficient. At the same time, compassion should not erase accountability. A restorative response can require the student to redo, explain, and learn from misrepresented work while distinguishing developmental failure from systematic deception.

Policies must also avoid linguistic injustice. Since detection systems can misclassify non-native English writing, professors should not use detector scores as proof. Oral conversation, comparison with earlier work, and direct questions about sources and decisions provide more educationally legitimate evidence.

Students using translation or language support should be assessed on the intended disciplinary capability, while remaining responsible for the accuracy and meaning of the final text.

7.10 Preserve human feedback, disagreement, and apprenticeship

Use AI to extend feedback capacity, not to withdraw faculty presence from the moments in which judgment and identity are formed. Automated comments may help with early practice or recurring errors. Professors and tutors should retain feedback on conceptual direction, professional readiness, serious failure, vocation, and work whose consequences matter. Peer review should remain a genuine encounter in which students learn to offer and receive correction.

In classrooms where students hesitate to speak because of language, hierarchy, or fear of losing face, AI can support preparation: students may rehearse a question, translate a contribution, or generate a counterargument before discussion. The teaching design should then invite that preparation into human dialogue. Professors can model respectful disagreement by revising their own claims, acknowledging uncertainty, and showing that correction need not destroy harmony.

7.11 Model the practices required of students

Professors should disclose material AI use in teaching and scholarship under the standards they expect from students. If AI assists in producing a case, feedback, translation, or course resource, the professor remains responsible for accuracy, bias, confidentiality, and pedagogical fitness. Students notice inconsistency when they are told to verify outputs while receiving generated materials that faculty have not checked.

Christian professors have an additional opportunity to model freedom from both anxiety and status competition. They can use powerful tools gratefully without presenting constant acceleration as a moral duty; acknowledge limits without loss of authority; credit collaborators and sources; and refuse uses that detach professional success from service. In secular institutions, these practices need not be announced as uniquely Christian, though their theological source can be stated transparently when relevant.

7.12 Test retention, transfer, and answerability

If an AI-enabled activity is intended to improve learning, inspect learning after the assistance changes or is removed. A short delayed question, unfamiliar application, oral explanation, code modification, error diagnosis, or comparison with a new case can reveal whether the student acquired a transferable capacity. Where AI use is itself the outcome, assess whether the student can specify the task, evaluate the output, recognize uncertainty, and decide when human or external evidence is required.

Professors do not need a randomized trial for every course change. They do need evidence better than enthusiasm, polished submissions, or end-of-course satisfaction. Compare cohorts where possible, inspect differential effects across levels of prior knowledge and language background, and invite

students to report both benefit and dependence. If output improves while independent explanation deteriorates, the course should treat that divergence as a design signal.

7.13 A practical course protocol

The recommendations can be condensed into a repeatable protocol.

Stage	Professor action	Student evidence
Before the course	Identify protected capabilities and AI-enabled professional capabilities	Clear learning outcomes and course-level AI statement
Before each major task	Classify use as required, permitted, staged, or prohibited; explain why	Task-specific acknowledgement of boundaries
During learning	Provide protected practice, constrained support, and opportunities for human correction	Initial attempts, questions, feedback responses, revisions
At submission	Request proportionate process and disclosure evidence	Source or decision trail, verification record, final ownership statement
At assurance points	Sample independent and AI-governed competence	Oral defence, unfamiliar problem, practical demonstration, or critique
After the task	Check retention, transfer, equity, and relational effects	Delayed performance and reflective account grounded in specific decisions
At course review	Revise the design in light of evidence and tool changes	Aggregated patterns rather than surveillance of individual students

This protocol does not eliminate professional judgment; it disciplines it. It asks professors to make explicit what good teaching has always required: what students are becoming capable of, which practices produce that development, how the instructor knows, and whether the means used are consistent with the ends pursued.

8. Recommendations for university leaders

AI governance is not principally the management of a new category of software. It is the stewardship of an institutional change that reaches curriculum, assessment, faculty work, student support, knowledge infrastructure, professional preparation, and the university's public credibility. A policy written only by an information-technology or academic-integrity office will therefore be too narrow. So will a strategy whose principal measures are licences purchased, users activated, hours saved, or ranking visibility gained.

University leaders should govern AI through three linked responsibilities. They must preserve the truthfulness of the university's claims about what its graduates know and can do; shape conditions in which students and staff can develop responsible agency; and direct institutional capabilities toward service rather than technological competition as an end in itself. These responsibilities require neither

centralized control over every classroom nor unrestricted local experimentation. They require a shared institutional architecture within which informed disciplinary judgment can operate.

8.1 Make educational purpose the governing layer

Every university AI strategy should begin with a short account of educational purpose that is capable of ruling out as well as authorizing uses. Statements such as “prepare AI-ready graduates,” “enhance productivity,” or “lead in digital transformation” are not sufficient. They identify capabilities or aspirations but do not specify the human and social goods those capabilities should serve.

A formation-sensitive strategy should state that AI adoption is subordinate to at least five commitments: truthful knowledge, intellectual agency, responsible character, vocational service, and human community. It should then translate these commitments into questions that can guide budgets and operating decisions. Will an advising chatbot extend access to a human adviser or become the institution's less costly substitute? Will automated feedback create more occasions for learning or enable the withdrawal of academic presence? Will research assistance help scholars inspect a wider body of evidence or normalize claims whose provenance no one has checked? Will administrative savings be returned to education and relationship, or absorbed into greater throughput?

The statement should also name what the institution will not optimize. A university should not maximize AI use for its own sake, treat reduced labor as an unqualified good, or regard every form of friction as waste. It should protect time for slow reading, foundational practice, laboratory and field experience, mentoring, collegial judgment, and encounter with persons affected by professional decisions. By making these commitments explicit, leaders give schools and professors stable reasons for changing specific rules as tools evolve.

For a Christian university, this purpose should be stated in theological language rather than inferred from a generic ethics policy. Human dignity is received rather than earned; knowledge is accountable to truth; intellectual gifts are vocations of service; persons are finite and fallen; and institutional hope does not rest in technological mastery. These claims should govern resource allocation, staffing, procurement, and assessment—not merely appear in a preamble. A university cannot credibly claim that education is relational while using AI savings primarily to reduce contact between faculty and students.

8.2 Require programme-level maps of capability and assurance

Course policies alone cannot preserve the integrity of a degree. A student may encounter many well-designed assignments yet graduate without a coherent demonstration of the programme's essential capabilities. Every academic programme should therefore produce a concise **capability and assurance map** containing four elements:

44. the disciplinary and professional capabilities graduates must possess;
45. the foundational components that must be learned or periodically demonstrated without generative assistance;
46. the AI-mediated capabilities expected in contemporary practice; and

47. the points at which the programme obtains trustworthy evidence of both.

This is not a call to restore an entirely tool-free curriculum. In software engineering, for example, a programme may require graduates to work effectively with coding assistants while retaining the ability to reason about algorithms, security, tests, architecture, and failure. A business programme may require AI-supported analysis while preserving independent competence in framing decisions, testing assumptions, interpreting evidence, and accounting for stakeholders. The map makes visible where automation is educationally appropriate and where it would undermine the claim represented by the qualification.

Assessment should then be designed across the programme rather than repaired assignment by assignment. TEQSA's assessment-reform guidance similarly argues that trustworthy judgments about learning require multiple, inclusive, and contextualized forms of evidence and meaningful assurance points across a programme. Formation-sensitive augmentation adds that assessment is not only an evidentiary device: it trains habits. A programme that responds to AI only through surveillance and invigilation may obtain some valid evidence while cultivating anxiety and adversarial relationships. One that accepts every polished artifact may cultivate misrepresentation and dependence. A balanced architecture combines assistance-rich learning, authentic AI-enabled professional work, process evidence, human dialogue, and a limited number of secure or live demonstrations of retained competence.

Programme review should ask not merely whether AI policy was followed, but whether students can explain unfamiliar problems, transfer knowledge, challenge generated outputs, use authoritative sources, communicate uncertainty, and take responsibility for consequences. These outcomes should be sampled across different levels of prior knowledge, language background, disability, and access. Aggregate evidence of a performance–learning divergence should trigger curriculum redesign even if pass rates and student satisfaction have risen.

8.3 Build distributed governance with clear authority

AI decisions cross domains of expertise. Academic leaders understand programme purposes; professors understand disciplinary learning; information-security and legal teams understand data and contractual risk; librarians understand provenance and knowledge infrastructure; disability and student-support specialists understand access; students experience the practices being governed; and affected professional communities understand downstream consequences. No single group should be permitted to define institutional benefit on behalf of all the others.

Universities should create a standing cross-disciplinary AI and education council with responsibility for principles, baseline safeguards, procurement review, evaluation standards, and appeals. Its membership should include academic staff from different disciplines, educational developers, students, librarians, technologists, privacy and security specialists, disability or accessibility expertise, professional staff, and external practitioners or community representatives where consequential deployments are contemplated. Christian institutions should include theological and pastoral expertise, but should not

assume that a theologian alone represents mission or that technical members are responsible only for implementation.

Authority should be distributed according to the level of decision:

Level	Primary responsibility	Institutional requirement
University	Mission, common safeguards, procurement, data governance, infrastructure, appeal	Publish principles, minimum controls, supported tools, decision rights, and review dates
Programme	Graduate capabilities, assessment assurance, professional standards	Maintain capability and assurance map; review aggregate learning evidence
Course	Task purpose, learner stage, permitted assistance, disclosure	Provide course- and task-specific classifications with educational rationale
Individual	Verification, attribution, judgment, responsible use	Remain answerable for claims and decisions; disclose material assistance

This distribution protects faculty judgment without turning academic freedom into an exemption from privacy, security, accessibility, or assessment-validity obligations. It also prevents central leaders from imposing a technically uniform practice on disciplines whose epistemic methods differ. Decisions should be contestable. Students and faculty need a clear route to question an automated classification, report harm, request an accommodation or non-AI route, and challenge a pedagogically incoherent requirement.

Governance should assume that silence is not consent. This is especially important in hierarchical institutions and in settings where students or junior staff may hesitate to criticize a senior leader's initiative. Anonymous feedback can help, but structured dissent is stronger: appoint a red-team reviewer for major pilots, require a recorded response to minority concerns, include junior members with voting rights, and review distributional effects before expansion. Such measures distinguish genuine harmony from compliance.

8.4 Treat procurement as an educational and moral decision

Licensing an AI system does more than purchase functionality. It can transfer student work, faculty materials, research data, decision records, and institutional dependence into a vendor ecosystem. It can also normalize a particular model of tutoring, advising, authorship, or evaluation before the university has examined whether that model serves education. Procurement must therefore test pedagogical fitness alongside price, security, and convenience.

At minimum, a procurement review should examine:

- the educational purpose and evidence supporting the proposed use;
- what data enter the system, where they are processed, how long they are retained, and whether they are used for model training;
- the availability of age-, disability-, language-, and context-appropriate safeguards;
- performance and failure patterns across relevant disciplines, languages, and student populations;

- source provenance, uncertainty communication, logging, auditability, and the ability of qualified humans to override outputs;
- ownership and reuse of prompts, teaching materials, student work, and generated content;
- integration costs, pricing changes, portability, model substitution, and a viable exit plan;
- labor effects, including whether promised efficiencies strengthen human educational work or principally remove it; and
- procedures for incident reporting, remedy, independent review, and vendor accountability.

Human review must be meaningful rather than ceremonial. A person cannot provide oversight without relevant expertise, adequate time, access to the information needed to evaluate the output, and authority to reject it. Where a system influences admission, progression, discipline, advising, disability support, or other consequential decisions, affected persons should know that automation is involved and have access to a timely human reconsideration. High-stakes student decisions should not rest on an AI-generated inference that no accountable person can explain.

Universities should be particularly cautious about tools that infer emotion, motivation, deception, engagement, or future success from behavioral traces. Such systems can convert pedagogical relationships into surveillance and reduce complex persons to contestable proxies. The *imago Dei* does not yield a technical rule for every prediction, but it establishes a strong presumption against classifications that are opaque, identity-defining, or difficult to contest.

An Asian perspective adds questions of technological sovereignty and localization. Institutions may reasonably seek models and infrastructure responsive to local languages, laws, and knowledge. Yet a locally hosted or nationally developed model is not thereby educationally or morally adequate. It may reduce one dependency while increasing another, encode local exclusions, or make institutional challenge more difficult. Leaders should test Mandarin and other Asian-language performance at the level of disciplinary meaning, not general fluency, and involve locally situated scholars and communities in evaluation.

8.5 Invest in faculty vocation, capability, and time

No policy can compensate for unsupported professors. Faculty development must join technical literacy, learning design, assessment, disciplinary epistemology, ethics, and practical experimentation. A short demonstration of prompting techniques may increase use without increasing pedagogical judgment. Professors need opportunities to redesign a real course, examine student work, compare assisted and unassisted performance, learn the failure modes of their discipline, and develop proportionate disclosure and verification practices.

Universities should recognize this work in workload and promotion systems. Redesigning assessments, testing tools, producing accessible alternatives, and evaluating learning take time. If leaders announce transformation while adding this labor invisibly to already stretched faculty, adoption will favor the most resourced departments and the most enthusiastic individuals. It may also undermine the very relationships the strategy claims to enhance.

Productivity gains require an explicit covenant. When AI reduces administrative work, leaders should state how the recovered capacity will be used. Returning time to feedback, mentoring, research, collegial deliberation, or rest is consistent with formation-sensitive augmentation. Simply increasing class size, output targets, or response expectations can convert efficiency into intensification. A Christian institution should be especially alert to this temptation. Sabbath is contradicted when every saved hour becomes a new institutional demand.

Faculty should also receive protection from compelled use that lacks pedagogical justification. A university may require baseline AI competence and common safeguards, but should permit a professor to delay or prohibit use in a particular activity when the target capability requires protected practice. Conversely, professional relevance may justify requiring students to use AI in an advanced course. Institutional policy should support both judgments when they are reasoned, transparent, accessible, and aligned with programme outcomes.

8.6 Govern for equity, linguistic plurality, and cultural context

Access is not achieved merely by providing every student an account. Students differ in devices, connectivity, paid features, prior knowledge, disability, confidence, language, privacy exposure, and access to human help. The same system may reduce a barrier for one learner and create dependence or disadvantage for another. Leaders should therefore measure who benefits, who is misclassified, who must disclose more, whose language is distorted, and who loses access to a human alternative.

Several safeguards follow:

- provide institutionally supported access when AI use is required, without assuming that a free consumer tier is equivalent;
- maintain accessible and pedagogically credible alternatives when students cannot use a tool for disability, privacy, religious, legal, or technical reasons;
- prohibit the use of AI-detection scores as sole evidence of misconduct, particularly given documented bias against non-native English writers;
- test translation and generation across the languages actually used by students and communities;
- audit whether local, minority, or non-Western sources are absent or systematically reframed;
- include students with disabilities, multilingual students, international students, and lower-income students in pilot evaluation; and
- prevent disclosure requirements from becoming excessive surveillance or a special burden on those already subject to suspicion.

Asian contextualization must extend beyond a generic appeal to Confucian culture. South Asia, Southeast Asia, East Asia, and diasporic settings differ in religious plurality, colonial history, language, regulation, infrastructure, and the authority accorded to family, teacher, institution, and state. Even within Confucian-heritage settings, mainland China, Singapore, Hong Kong, and Taiwan cannot be treated as variations of one educational culture. UNESCO's Asia-Pacific work appropriately warns of uneven readiness and cultural representation. Institutions should therefore conduct local inquiry rather than infer student needs from national stereotypes.

At the same time, the region offers intellectual resources for governance. Traditions of disciplined learning challenge instant mastery; relational accounts of personhood challenge individualized content delivery; filial and social responsibility challenge purely private models of career success; and harmony, properly understood, can sustain difference within relationship. These resources should be received critically. Discipline can become overwork, relationality can become coercion, and harmony can become silence. Formation-sensitive governance asks which version a particular AI-mediated practice is strengthening.

8.7 Evaluate pilots as educational interventions

An institutional AI pilot should be treated as a revisable educational intervention, not as the first stage of inevitable rollout. Before launch, leaders should record the purpose, comparison or baseline, target population, expected mechanism, material risks, outcome measures, stopping conditions, review date, and owner of the final decision. This makes it possible to learn from a failed pilot without treating failure as institutional embarrassment.

Evaluation should include four layers:

48. **Implementation:** Who used the system, how, and with what support? Did actual practice match the design?
49. **Immediate effects:** Did access, feedback, completion, quality, or workload change?
50. **Formation and learning:** Did students retain and transfer the target capability, verify more effectively, exercise better judgment, disclose honestly, and continue to seek human correction?
51. **Distribution and system effects:** Which groups benefited or were burdened? Were faculty roles, relationships, costs, data practices, or vendor dependence altered?

Adoption, satisfaction, and output quality belong mainly to the first two layers. They cannot establish the latter two. Where rigorous causal evaluation is feasible, universities should use comparison groups, delayed measures, and preregistered analyses. Where it is not, a mixed-method design can combine performance tasks, oral checks, process artifacts, interviews, focus groups, usage records, and equity analysis. Data collection should remain proportionate; the desire to study formation must not justify pervasive monitoring.

The final decision should allow four outcomes: expand, revise, restrict, or discontinue. A pilot that saves time but weakens learning may be redesigned so that AI supplies hints rather than answers. A system that helps multilingual students but performs poorly with local sources may be retained for language support and prohibited for evidence synthesis. Reversibility should be part of the original contract and technical design rather than discovered after dependence has formed.

8.8 Give Christian institutions a mission-specific account

Christian universities face a distinctive temptation: to adopt the same systems and metrics as their competitors while adding devotional language to the strategy. Mission alignment requires more. Governing bodies and senior leaders should be able to explain how specific investments reflect creation,

creatureliness, fallibility, truth, vocation, community, grace, and hope. They should also invite scrutiny of whether institutional practices contradict those claims.

Five questions can structure a mission review:

52. **Dignity:** Does the practice attend to each student as a person, or primarily as a risk score, productivity unit, customer, or source of data?
53. **Truth:** Does it strengthen warranted belief, honest authorship, and institutional credibility, or reward plausible appearance?
54. **Creatureliness:** Does it remove needless burden while preserving formative effort, rest, embodiment, and dependence?
55. **Vocation and neighbor:** Whose good is served, who bears residual risk, and can responsible persons answer to those affected?
56. **Community and hope:** Does it return learners to communities of correction and care, and is the institution free to refuse technological promises that conflict with its mission?

The review should include students, faculty, staff, governing-board members, churches or ecclesial partners where appropriate, and communities affected by professional education. Christian identity does not immunize the institution from market anxiety, prestige competition, or paternalism. Christian accounts of fall and continuing reform give Christian universities reason to welcome external critique and publish evidence even when it complicates a success narrative.

The university should also develop practices of worship, pastoral care, and communal reflection that address AI-mediated anxiety and desire. Policy cannot by itself loosen the connection between achievement and worth. Students need communities that embody grace: places where assistance can be sought without shame, failure can be confessed and repaired, excellence is celebrated without becoming salvation, and vocation is discerned in relation to neighbor rather than ranking alone. These practices are not substitutes for rigorous teaching or governance. They are part of the formative ecology in which rules acquire meaning.

8.9 Equip Christian leaders in plural universities for public faithfulness

Christian professors and leaders in secular universities should neither conceal the sources of their convictions nor presume that colleagues share them. They can work in two registers. Theologically, they may understand human dignity through creation, limits through creatureliness, institutional risk through the fall, professional responsibility through vocation, and confidence without technological anxiety through grace and hope. Publicly, they can offer reasons accessible to colleagues: reliable learning, valid credentials, truthful authorship, professional accountability, human rights, equitable access, cultural plurality, meaningful consent, and the common good.

This translation is not a dilution of faith. It is an act of neighborly reasoning within an institution composed of people with different ultimate commitments. Christians should seek coalitions with colleagues who share particular educational goods, receive correction from other traditions and disciplines, and resist making AI governance a proxy contest over religious identity. They can contribute

a distinctive warning against both reduction of persons to performance and the hope that technique will resolve moral and institutional problems without costly human change.

Concrete leadership may be modest: asking what capability a proposal forms, insisting on an exit clause, protecting a human adviser, requesting evidence after assistance is removed, including a junior student voice, or refusing an unsupported claim of inevitability. Such acts embody faithful presence more credibly than a sweeping declaration for or against AI.

8.10 A staged institutional roadmap

The recommendations can be organized into a sequence that permits immediate action without pretending that the university already possesses final answers.

Horizon	Priority actions	Evidence of progress
First six months	State educational and mission principles; establish decision rights and a cross-disciplinary council; publish baseline privacy, security, disclosure, and appeal safeguards; inventory current uses	Public principles; named accountable owners; risk register; supported-tool list; accessible reporting route
Six to eighteen months	Produce programme capability maps; identify assurance checkpoints; fund faculty redesign cohorts; review major contracts; launch a small portfolio of evaluable pilots	Programme maps; redesigned assessments; faculty workload allocation; procurement records; pilot protocols with baselines and stopping conditions
Eighteen to thirty-six months	Compare retention, transfer, equity, and relational outcomes; revise policy by evidence; strengthen local-language evaluation; address labor and vendor-dependence effects	Evaluation reports; documented expansion, redesign, restriction, or termination decisions; equity and language audits; exit tests
Continuing	Review annually and after material model or regulatory changes; renew theological or mission discernment; publish learning and unresolved risks	Versioned policy; transparent incident and evaluation summaries; sustained human educational contact; research partnerships

The roadmap is intentionally iterative. Current public guidance converges on human-centered use, teacher capability, assessment reform, and learning assurance (Miao and Holmes 2023; Miao and Cukurova 2024; OECD 2026; TEQSA 2025). Formation-sensitive augmentation supplies the integrating question: do these institutional arrangements cultivate and truthfully certify people who can understand, judge, relate, and remain answerable in an AI-mediated world?

9. Research agenda: from performance effects to human formation

The present evidence base can inform near-term design, but it cannot yet support confident claims about the long-run formation of students or institutions. Most empirical studies examine a course, a small set of tasks, or a brief period of use. Common outcomes include immediate performance, motivation, satisfaction, self-reported effort, or attitudes. These are valuable but incomplete. University

leaders need to know whether capabilities persist, transfer to unfamiliar settings, support responsible professional judgment, and develop differently across student populations and institutional cultures.

Theological argument cannot fill these empirical gaps. The doctrines developed in Section 4 identify goods, risks, and questions; they do not prove that a particular chatbot causes passivity, strengthens pride, or weakens community. Conversely, empirical research can measure behavior and outcomes without deciding whether efficiency, autonomy, harmony, or employability should be ultimate educational aims. A rigorous research programme must keep these tasks distinct while bringing them into disciplined conversation.

9.1 Longitudinal learning and the performance–learning divergence

The highest priority is to distinguish assisted performance from durable capability. Research should follow students beyond the task on which AI was used and examine delayed retention, transfer to unfamiliar problems, ability to detect error, quality of explanation, and performance when the form of assistance changes. Studies should compare at least three conditions where feasible: protected human practice, pedagogically constrained AI support, and unrestricted general-purpose assistance.

The key questions are developmental rather than merely comparative. At what point does assistance enable a learner to attempt more complex work, and at what point does it bypass construction of the mental model needed to evaluate the answer? Does a sequence of unaided attempt, constrained feedback, and later general-purpose use outperform permanent access or permanent prohibition? How much independent competence must be retained in a field where tools are normally available, and how often should it be sampled?

These questions require semester- and programme-level studies, not only single-session experiments. They also require attention to attrition and compensatory behavior: students assigned to a restricted condition may use unsanctioned tools, while students assigned AI may differ in how much they delegate. Process evidence and interviews are therefore needed to interpret outcome measures.

9.2 System design and the distribution of cognitive work

“Use of generative AI” is too broad an intervention category. A tutor that asks questions and releases graduated hints differs educationally from a system that generates a final solution; a retrieval-grounded tool differs from a fluent model with no visible sources; a discipline-specific system differs from a consumer chatbot. Future studies should report the model and version, interface, system instructions, grounding sources, memory, feedback design, permitted actions, and exact pattern of learner interaction.

Research should manipulate how cognitive work is distributed. Useful comparisons include answer-first versus question-first interfaces; immediate solutions versus hints; learner explanation before versus after assistance; visible sources versus unsupported claims; generated critique versus generated content; persistent personalization versus session-limited support; and systems optimized for agreement versus systems designed to surface uncertainty and counterargument.

The outcome should not be reduced to whether students like the system. Sycophantic assistance may be preferred precisely because it reduces the discomfort of correction. Studies should examine calibration of confidence, willingness to revise, source checking, help-seeking, and the capacity to state what remains unknown. System design can then be evaluated as pedagogy rather than user experience alone.

9.3 Authorship, truthfulness, and moral agency

Academic-integrity research has concentrated heavily on detection and rule violation. A formation-sensitive agenda asks how different assessment and disclosure regimes shape truthful practice. Do clear, task-specific permissions increase honest disclosure? Do burdensome prompt logs produce ritual compliance or greater understanding? Does an oral defence promote ownership, or merely add anxiety? When a student misrepresents AI-supported work, which responses best restore learning and trust?

These questions should be studied without pretending that virtue can be read directly from behavioral traces. Researchers can observe disclosure accuracy, source verification, willingness to correct a claim, acceptance of responsibility, and decisions under temptation. They can interview students about shame, pressure, uncertainty, and the meaning of authorship. These are indicators of practices relevant to character, not complete measurements of an inward virtue.

Asian and Chinese contexts are especially important here. Researchers should examine how face, family expectations, high-stakes competition, teacher authority, and linguistic insecurity affect disclosure and help-seeking. The hypothesis should not be that “Asian students” are more conformist or instrumental. Comparative designs should test how particular institutional and relational conditions operate within and across settings. Reciprocal family support may strengthen agency, whereas shame-based pressure may increase concealment; respect for teachers may support apprenticeship or inhibit correction depending on how authority is exercised.

9.4 Relational learning, presence, and correction

AI tutoring research usually measures learning and engagement, but rarely whether the system changes student relationships with teachers and peers. Does an always-available assistant enable students to prepare better questions and participate more fully, or does it reduce office-hour use and peer discussion? Which students experience AI as a bridge into community, and which use it to avoid vulnerability? Do students become less willing to tolerate imperfect, slower, or disagreeable human partners after extensive interaction with personalized systems?

Studies should combine network measures, observation, interviews, and learning outcomes. The number of messages sent to a chatbot is not evidence of relationship, and frequency of classroom interaction is not evidence of meaningful belonging. Researchers should inspect whether students give and receive correction, seek help from appropriate persons, contribute to shared inquiry, and sustain disagreement without withdrawal.

There is also a need to distinguish educational tutoring from companionship or pastoral simulation. Systems that use empathic language may provide immediate support, particularly where human services

are inaccessible. Yet universities should study over-reliance, displaced human help, crisis escalation, privacy, and the expectations students form about reciprocity. The moral significance of simulated care cannot be inferred from user satisfaction alone.

9.5 Vocational judgment and consequential practice

Universities do not only prepare students to complete academic tasks. They prepare engineers, clinicians, teachers, managers, lawyers, researchers, public servants, designers, and computing professionals whose decisions affect others. Research should therefore move into authentic and simulated professional settings where error, uncertainty, trade-offs, and accountability are visible.

Key outcomes include whether learners can identify missing stakeholders, challenge an AI recommendation, retrieve authoritative evidence, communicate uncertainty, recognize distribution shift, document rejected alternatives, and explain a final decision to an affected person. Studies should compare students who learn with AI throughout, those who acquire foundations before AI use, and experienced professionals. This can reveal whether AI changes the threshold expertise required for safe practice and which forms of independent competence remain essential.

Participatory research is particularly important. Patients, clients, workers, community partners, and minority-language groups should help define what counts as responsible performance. A technically accurate recommendation may still be poorly fitted to local practice, difficult to contest, or damaging to a relationship. Those who bear residual risk should not enter the study only as outcome data.

9.6 Faculty vocation and institutional transformation

The formative effects of AI are not confined to students. Studies should examine how automated feedback, content generation, advising, research assistance, and administrative automation alter faculty attention, expertise, authority, workload, and professional identity. Does saved time return to mentoring and inquiry, or is it converted into larger classes and higher output expectations? Does reliance on generated teaching material weaken disciplinary freshness or enable more responsive education? Which academic roles are devalued, and who performs the new labor of checking, correcting, and caring for exceptions?

Institutional research should analyze governance and political economy. Relevant variables include vendor contracts, data flows, centralized versus distributed authority, faculty participation, appeal mechanisms, local-language performance, cost escalation, and the ability to exit. Comparative case studies could examine why similar tools become augmentative in one institution and substitutive in another. The unit of analysis should include the organization and its incentives, not only the individual user.

Christian universities offer a particularly important but underexamined setting. Research should test whether a stated theology of personhood changes actual procurement, staffing, assessment, and student-support decisions. Mission statements may discipline adoption, or they may function as

legitimating language around strategies driven by the same competitive pressures as elsewhere. Constructive institutional theology requires evidence about practice.

9.7 Asia, Chinese contexts, and comparative contextual research

Much of the visible literature and many benchmark assumptions remain Anglophone. A serious Asian research agenda should not simply replicate Western studies with a different sample. It should investigate how multilingual reasoning, local knowledge, family obligation, examination systems, educational hierarchy, religion, infrastructure, and national development strategies alter both AI use and the meaning of educational outcomes.

Priority studies include:

- comparisons of AI-supported reasoning in English, Mandarin, and other Asian languages using discipline-specific sources;
- effects of translation assistance on conceptual precision, voice, participation, and assessment fairness;
- student and faculty practices across mainland China, Singapore, Hong Kong, Taiwan, and other Asian systems without presuming cultural equivalence;
- interaction between AI use, private tutoring markets, socioeconomic inequality, and differential access to premium systems;
- the influence of teacher authority and face on error correction, disclosure, and willingness to challenge confident outputs;
- the relationship between national AI-literacy initiatives and universities' capacity for independent educational judgment; and
- comparative theological work involving Protestant, Catholic, Orthodox, Confucian, Buddhist, Muslim, Hindu, and secular humanist accounts of learning and personhood.

The goal is neither to discover an essential “Asian response” nor to make Western categories universal. It is to identify which mechanisms travel, which goods are differently ordered, and how local communities should participate in the design and judgment of educational systems.

9.8 A common research architecture

The field would benefit from a shared minimum architecture that preserves comparability without suppressing disciplinary difference.

Research element	Minimum expectation
Context	Institution, country or region, discipline, course level, language, stakes, and mode of study
Learners	Prior knowledge, relevant demographic and linguistic variables, access conditions, and attrition
AI intervention	Model and version, interface, system design, grounding, allowed functions, duration, and actual use
Comparator	Business-as-usual, protected practice, constrained AI, or another clearly specified condition

Research element	Minimum expectation
Immediate outcomes	Task performance, time, engagement, workload, access, and error patterns
Durable outcomes	Delayed retention, transfer, explanation, calibration, independent performance, and governed AI use
Formative practices	Verification, attribution, revision, help-seeking, correction, responsibility, and human participation
Distribution	Differential benefit and harm by prior knowledge, language, disability, socioeconomic access, and other relevant groups
Institutional effects	Faculty labor, human contact, costs, data governance, vendor dependence, and appeal
Transparency	Preregistration where appropriate, materials and prompts, deviations, funding, conflicts, limitations, and model changes

Mixed methods are normally required. Quantitative measures can identify changes in performance, retention, calibration, and distribution. Qualitative work can reveal why learners delegate, what they believe they own, how authority and face operate, and whether an apparently efficient practice changes the meaning of teaching or vocation. Neither should be asked to do the other's work.

The rapidly changing model environment also creates a replication problem. Researchers should archive prompts and interface conditions where legally possible, record dates and model identifiers, and conduct sensitivity checks after material updates. Findings should be framed around mechanisms—answer provision, hinting, source visibility, personalization, sycophancy—rather than treated as timeless properties of a brand name.

9.9 A constructive programme for Christian scholarship

Christian scholarship can contribute at four levels. First, theology clarifies the goods at stake and exposes hidden anthropologies within ostensibly neutral measures. Second, philosophy and educational theory translate concepts such as agency, authorship, wisdom, vocation, and community into contestable accounts of educational practice. Third, empirical research tests proposed mechanisms and attends to students, faculty, and affected communities. Fourth, institutional action embodies and revises the framework through curriculum, policy, budgets, and relationships.

This programme should resist two shortcuts. It should not baptize generic “responsible AI” principles by attaching biblical vocabulary, and it should not treat a plausible theological concern as if it were an observed effect. Claims about idolatry, dignity, vocation, and grace are theological judgments; claims about retention, dependence, disclosure, or relational displacement require empirical evidence. Their interaction is the research contribution: theology helps decide what to investigate and why it matters, while evidence disciplines claims about what a technology is actually doing.

A practical next step would be a multi-institutional consortium of Christian and secular universities in Asia. Participating programmes could share a common set of delayed-learning and answerability measures while conducting discipline-specific pilots. Qualitative work could examine family obligation,

face, language, teacher authority, vocation, and institutional mission. Christian institutions could compare whether explicit theological commitments generate measurably different practices, not simply different rhetoric. Such a consortium would move the debate from generalized hope and fear toward accountable, contextually grounded learning.

10. Conclusion: faithful presence without fear or idolatry

Generative AI has made visible a question universities could previously postpone: when an academically acceptable performance can be produced without the learner possessing the corresponding understanding, what exactly is the university trying to form and certify? No policy about permitted tools can answer that question. It requires an account of education, the human person, and the goods toward which knowledge should be directed.

The empirical evidence gives universities neither permission for complacency nor grounds for panic. Structured AI tutors and well-designed assistance can widen access, provide rapid feedback, and produce meaningful near-term learning gains. Unrestricted assistance can also improve performance while leaving weaker independent capability. Research on cognitive offloading, ownership, critical thinking, and sycophancy identifies plausible mechanisms of deformation, but the evidence remains young, context-dependent, and short-term. The honest conclusion is conditional: AI can support learning when it elicits thought, preserves epistemic control, and operates within structures of human responsibility; it can undermine learning when it supplies the appearance of competence while bypassing the practices through which competence and judgment grow.

Evangelical Christian theology deepens the issue without supplying a simplistic verdict on the technology. Creation affirms human making and the unearned dignity of every learner. Creatureliness teaches that limits, dependence, effort, embodiment, and rest are not all defects. The fall directs suspicion toward students, faculty, leaders, institutions, vendors, and cultural ambitions alike. Incarnation and community affirm forms of knowing that require presence, participation, correction, and care. Truth and wisdom distinguish fluent language from responsible judgment. Vocation orders capability toward answerable service. Grace frees students and institutions from making achievement or innovation the source of worth. Hope frees leaders to act energetically without imagining that technological mastery will redeem education or secure the future.

The dialogue with Chinese and wider Asian contexts sharpens each claim. Reverence for learning, disciplined practice, relational responsibility, family obligation, and harmony provide real resources against shallow individualism and frictionless automation. They also have distorted forms: credentialism, overwork, authoritarian obligation, deference, and conformity. AI may reduce linguistic barriers and fear of losing face, yet also intensify competition, privatize learning, amplify confident authority, and marginalize local knowledge. Contextual faithfulness therefore means neither defending a static culture nor importing a Western template. It means discerning which goods and distortions a practice strengthens in a particular community.

Formation-sensitive augmentation is the paper's practical response. It begins with the purpose of an educational activity, identifies the capacity and relationship being formed, distinguishes delegated execution from human answerability, and varies assistance according to learner development, system design, and stakes. It asks universities to require AI where governed use is itself a necessary capability, permit it where it does not displace the educational purpose, stage or restrict it where novices require protected practice, and prohibit it narrowly where valid certification, safety, privacy, or non-delegable relationship demands this.

The framework does not preserve a romantic image of education untouched by tools. Human inquiry has always depended on languages, books, instruments, institutions, teachers, and inherited knowledge. Nor does it define human distinctiveness as the shrinking residue of tasks machines cannot perform. The issue is not whether an AI executed part of the work. It is whether learners and accountable communities remain able to understand what matters, test what is claimed, judge what should be done, explain their decisions, receive correction, and answer to those who bear the consequences.

For professors, this means protecting some novice practices, staging assistance, teaching discipline-specific AI literacy, assessing process where process matters, preserving human feedback, and checking retention and transfer. For university leaders, it means grounding strategy in educational purpose, mapping programme capabilities, reforming assessment for learning assurance, governing procurement as an educational decision, investing in faculty time, protecting linguistic and material equity, and evaluating pilots beyond adoption and satisfaction. For Christian institutions, it also means allowing theology to govern budgets, staffing, data, relationships, and institutional ambition—not merely the language of policy. For Christians serving in plural universities, it means offering these convictions through transparent public reasoning and patient coalition-building.

Universities will make mistakes. Models will change, evidence will be revised, and practices that appear beneficial may generate unexpected dependencies or exclusions. A Christian response should therefore be marked by institutional humility: readiness to test, confess, repair, and change course. Reversibility is not weakness; it is an acknowledgement of fallibility. Refusing inevitability is not fear; it is the exercise of responsible judgment.

The decisive measure of an AI strategy is consequently not how extensively a university automates, how rapidly it adopts, or how convincingly its students perform with assistance. It is whether the university forms and truthfully certifies persons capable of seeking truth, exercising wise and courageous judgment, participating in communities of correction and care, accepting responsibility for AI-supported work, and directing knowledge toward the service of God and neighbor. Universities should welcome every use of AI that serves that calling and refuse every use that asks them to abandon it.

References

Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman. 2025. “Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics.” *Proceedings of the National Academy of Sciences* 122 (26): e2422633122. https://doi.org/10.1073/pnas.2422633122.

Calvin, John. 1559/1960. Institutes of the Christian Religion. Translated by Ford Lewis Battles, edited by John T. McNeill. 2 vols. Philadelphia: Westminster Press.

Chan, Simon. 1998. Spiritual Theology: A Systematic Study of the Christian Life. Downers Grove, IL: InterVarsity Press.

Chen, Shuzhen, and Alan C. K. Cheung. 2025. "Effect of Generative Artificial Intelligence on University Students' Learning Outcomes: A Systematic Review and Meta-Analysis." *Educational Research Review* 49: 100737. https://doi.org/10.1016/j.edurev.2025.100737.

Cheng, Myra, Cinoo Lee, Pranav Khadpe, Sunny Yu, Dyllan Han, and Dan Jurafsky. 2026. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." *Science* 391 (6792): eaec8352. https://doi.org/10.1126/science.aec8352.

Deng, Ruiqi, Maoli Jiang, Xinlu Yu, Yuyan Lu, and Shasha Liu. 2025. "Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies." *Computers & Education* 227: 105224. https://doi.org/10.1016/j.compedu.2024.105224.

Digital Education Council. 2026. *AI in Higher Education: Global Survey 2026*. Report page.

Chow, Alexander, and Stephanie Wong. 2023. "Chinese Christian Theology." *St Andrews Encyclopaedia of Theology*. https://www.saet.ac.uk/Christianity/ChineseChristianTheology.

Hansen, Traynor. 2025. "Faithful Writing Pedagogy in the Age of Generative AI: A Sabbath-Grounded Approach." *Christian Scholar's Review* 54 (2): 55–76. Article.

Holmes, Arthur F. 1987. *The Idea of a Christian College*. Rev. ed. Grand Rapids, MI: Eerdmans. Publisher record.

Kestin, Greg, Kelly Miller, Anna Klales, Timothy Milbourne, and Gregorio Ponti. 2025. "AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting." *Scientific Reports* 15: 17458. https://doi.org/10.1038/s41598-025-97652-6.

Kuyper, Abraham. 1880/2015. "Sphere Sovereignty." In Abraham Kuyper: A Centennial Reader, edited by James D. Bratt, 461–490. Grand Rapids, MI: Eerdmans.

Kuyper, Abraham. 1898/1931. Lectures on Calvinism: The Stone Lectures Delivered at Princeton University. Grand Rapids, MI: Eerdmans.

Lai, Pan-Chiu. 2003. "Inheriting Chinese and Christian Traditions in Global Context: A Confucian-Protestant Perspective." *Religion & Theology* 10 (1): 1–23. Publisher record.

Lam, Chi-Ming. 2026. "Confucian Self-Cultivation and the Ethics of Learning with Large Language Models." *Studies in Philosophy and Education*. Advance online publication. https://doi.org/10.1007/s11217-026-10045-8.

Langford, Michael D. 2022a. "A Theological Framework for Reflection on Artificial Intelligence." In *AI, Faith, and the Future: An Interdisciplinary Approach*, edited by Michael J. Paulus Jr. and Michael D. Langford, 70–94. Eugene, OR: Pickwick. Open-access record.

Langford, Michael D. 2022b. "Artificial Intelligence and Theological Personhood." In *AI, Faith, and the Future: An Interdisciplinary Approach*, edited by Michael J. Paulus Jr. and Michael D. Langford, 95–127. Eugene, OR: Pickwick. Book record.

Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. 2025. "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task." arXiv:2506.08872. MIT project and limitations.

Lee, Hao-Ping, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers." In *Proceedings of CHI 2025*. https://doi.org/10.1145/3706598.3713778.

Leo XIV. 2026. *Magnifica Humanitas*. Encyclical letter, 15 May 2026. Official text.

Li, Changyue, Hang Cui, and Linda Serra Hagedorn. 2026. "The Cognitive Impact of ChatGPT in Higher Education: A Systematic Review of Critical and Creative Thinking Outcomes." *Computers and Education: Artificial Intelligence* 10: 100571. https://doi.org/10.1016/j.caeai.2026.100571.

Liang, Weixin, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. 2023. "GPT Detectors Are Biased against Non-Native English Writers." *Patterns* 4 (7): 100779. https://doi.org/10.1016/j.patter.2023.100779.

McGrath, Alister E. 2005. Iustitia Dei: A History of the Christian Doctrine of Justification. 3rd ed. Cambridge: Cambridge University Press.

OECD. 2026. *OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education*. Paris: OECD Publishing. https://doi.org/10.1787/062a7394-en.

Smith, David I. 2018. *On Christian Teaching: Practicing Faith in the Classroom*. Grand Rapids, MI: Eerdmans. Publisher record.

Smith, James K. A. 2009. *Desiring the Kingdom: Worship, Worldview, and Cultural Formation*. Grand Rapids, MI: Baker Academic. Publisher record.

State Council of the People's Republic of China. 2026. "China Aims to Build an AI Literacy System." 15 April 2026. Official English-language report.

Stephenson, Rose, and Charlotte Armstrong. 2026. *Student Generative Artificial Intelligence Survey 2026*. HEPI Report 199. Oxford: Higher Education Policy Institute. Report.

Volf, Miroslav, and Matthew Croasmun. 2019. *For the Life of the World: Theology That Makes a Difference*. Grand Rapids, MI: Brazos Press. Publisher record.

UNESCO. 2023. "How Generative AI Is Reshaping Education in Asia-Pacific." Updated 11 November 2025. https://www.unesco.org/en/articles/how-generative-ai-reshaping-education-asia-pacific.

Miao, Fengchun, and Mutlu Cukurova. 2024. *AI Competency Framework for Teachers*. Paris: UNESCO. Official publication.

Miao, Fengchun, and Wayne Holmes. 2023. *Guidance for Generative AI in Education and Research*. Paris: UNESCO. Updated 16 January 2026. Official publication.

Miao, Fengchun, Kelly Shiohira, and Natalie Lao. 2024. *AI Competency Framework for Students*. Paris: UNESCO. Official publication.

Tertiary Education Quality and Standards Agency (TEQSA). 2025. *Enacting Assessment Reform in a Time of Artificial Intelligence*. Melbourne: TEQSA. Official publication.

Wolterstorff, Nicholas. 2004. *Educating for Shalom: Essays on Christian Higher Education*. Edited by Clarence W. Joldersma and Gloria Goris Stronks. Grand Rapids, MI: Eerdmans. Publisher record.

Xia, Qi, Weijia Li, Yiming Yang, Xiaojing Weng, and Thomas K. F. Chiu. 2025. "A Systematic Review and Meta-Analysis of the Effectiveness of Generative Artificial Intelligence (GenAI) on Students' Motivation and Engagement." *Computers and Education: Artificial Intelligence* 9: 100455. https://doi.org/10.1016/j.caeai.2025.100455.

Yong, Amos. 2000. Discerning the Spirit(s): A Pentecostal-Charismatic Contribution to Christian Theology of Religions. Sheffield: Sheffield Academic Press.

Zhou, Jing, Qingke Guo, and Ruru Xu. 2020. "Reciprocal Filial Piety Facilitates Academic Success via Autonomy: Generalizing Findings in Chinese Society to a Global Context." *Frontiers in Psychology* 11: 69. https://doi.org/10.3389/fpsyg.2020.00069.