

THEOLOGY & PUBLIC LIFE WHITE PAPER

Remaining Human

Artificial Intelligence, Human Dignity, and Christian Hope

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

*“AI should extend human capacity without displacing truth-seeking,
moral responsibility, reciprocal relationships, embodied care, and
dependence on God”*

Hoong Chuin Lau

Updated 8 AUGUST 2026

For Christian leaders, professionals, educators, policymakers, technologists, and churches

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.

Content

Abstract

Executive summary

1. Scope, method, and the question before us

2. What AI is, and what it tempts us to believe

3. A Christian account of humanity

4. How AI may reshape human life

5. Singapore and Asia: contextual discernment

6. Seven tests for faithful AI stewardship

7. A practical agenda for Christian witness and public responsibility

8. A research agenda for Singapore and Asia

9. Conclusion: faithful construction between Babel and Jerusalem

Appendix A. Biblical map for AI and humanity

Appendix B. Ten questions for a board, ministry, or project team

References

Note on interpretation

Abstract

Artificial intelligence is often discussed as a contest between technological opportunity and technological risk. That framing is necessary but insufficient. AI is also an anthropological and spiritual question because it mediates how people know, work, communicate, create, govern, relate, and imagine their future. The decisive question is not whether machines will become sufficiently human, but whether human societies will adopt machine-like accounts of intelligence, value, freedom, and flourishing.

This white paper develops a Christian Christian framework for discerning AI's impact on humanity, with particular attention to Singapore and Asia. It places Scripture in dialogue with John C. Lennox's *2084 and the AI Revolution*, Pope Leo XIV's *Magnifica Humanitas*, selected Christian scholarship, empirical studies, and current Singapore and ASEAN policy. The paper argues that AI is neither an autonomous moral person nor a technological savior. It is a powerful family of human-made socio-technical systems: shaped by data, institutions, capital, political authority, design choices, and human desires. Its benefits can be substantial in science, healthcare, accessibility, public service, creativity, and work. Yet those same systems can erode truthful judgment, cultivate dependence, simulate intimacy without reciprocity, intensify surveillance and inequality, obscure responsibility, reduce work to measurable output, and reinforce dreams of mastery over human limitation and death.

The paper proposes seven tests for faithful AI stewardship: truth, dignity, agency, neighbor-love, justice, creatureliness, and hope. It applies them to personal use, families, churches, workplaces, technology development, and public governance. Its conclusion is neither rejection nor surrender. Christians are called to receive technical gifts gratefully, resist idolatrous promises, protect vulnerable neighbors, build accountable institutions, and bear hopeful witness to the human person as embodied image-bearer, responsible creature, sinner in need of grace, and participant in the resurrection life of Jesus Christ.

Executive summary

The central claim

AI is not merely a neutral instrument placed into an otherwise unchanged society. It is a *mediating environment*: it changes what is easy, visible, rewarded, measurable, and expected. Repeated reliance on AI can therefore reshape human habits and institutions even when no individual output is obviously harmful. The Christian task is to ask not only, “What can this system do?” and “Is it accurate?” but also, “What kind of person, relationship, workplace, church, and society does its use encourage us to become?”

The paper's thesis is:

AI is a powerful form of mediated human agency that may be used in service of love and the common good, but it must never be allowed to define human worth, absorb moral responsibility, imitate relationships so as to replace communion, or turn technical power into a promise of salvation.

Seven Christian convictions

1. **Human dignity is given, not earned.** Every person bears God's image (Gen 1:26-28; 9:6; Jas 3:9). Worth does not rise with cognitive performance, productivity, health, social status, or technological enhancement.
2. **Human beings are embodied and relational creatures.** We are not disembodied information processors. The creation of embodied life, the incarnation of the Word, and the resurrection of the body resist accounts of humanity that reduce persons to data or transferable patterns (Gen 2:7; John 1:14; 1 Cor 15).
3. **Intelligence is not wisdom.** Scripture locates wisdom in reverence for God, moral formation, truthful judgment, and faithful action—not in computational range or verbal fluency (Prov 1:7; 9:10; Jas 3:13-18).
4. **Power is corrupted by sin.** AI systems embody human incentives and institutional structures. The fall warns against both naïve technological optimism and fatalistic claims that innovation is uncontrollable (Gen 3; 11:1-9; Rom 1:21-25).
5. **Work is vocation and service, not merely output.** Automation can remove drudgery and enlarge service, but work also carries goods of contribution, skill, responsibility, solidarity, and provision for others (Gen 2:15; Eph 4:28; Col 3:23-24).
6. **Freedom is answerable agency, not frictionless choice.** Christian freedom is freedom for love and faithful obedience, not liberation from all dependence or limits (Gal 5:13; 1 Pet 2:16).

7. **Christian hope is resurrection, not technical immortality.** Healing and longevity are goods, but salvation is received through Christ rather than engineered through self-enhancement (John 11:25-26; Eph 2:8-10; 1 Cor 15:42-57; Rev 21:1-5).

What the evidence supports

- AI can increase productivity, accessibility, scientific discovery, and the quality of some task outputs. Exposure, however, does not by itself determine whether jobs will be displaced or transformed; design choices, bargaining power, training, and social policy matter.
- AI can improve individual creative performance while making outputs across people more similar. It can assist human creativity without necessarily enlarging a culture's collective imagination.
- Conversational AI can provide useful psychological support in designed interventions, but relational vulnerability also predicts stronger engagement. Benefits do not erase concerns about dependence, conflicts of interest, or substitution for accountable human care.
- Sycophantic models can strengthen users' belief in their own rightness and reduce willingness to repair relationships. An agreeable system may therefore be experienced as helpful while weakening moral correction.
- AI-generated and AI-mediated content can influence beliefs, amplify deception, and blur authorship. The societal problem is not only isolated falsehood but the erosion of shared practices for establishing truth.
- Environmental costs are material. The International Energy Agency projects that global data-centre electricity use will roughly double by 2030, with AI a major driver. Efficiency gains do not remove stewardship obligations.

These findings remain domain- and design-dependent. The literature is developing rapidly and rarely measures long-term moral, relational, spiritual, or civic formation. The paper therefore distinguishes established evidence, plausible mechanisms, theological judgment, and speculation.

Why Singapore and Asia matter

Singapore combines unusually high AI exposure, strong state capacity, global technology investment, a highly educated workforce, and a multicultural, multireligious population. The country's updated National AI Strategy continues to frame AI as serving the public good, while Singapore and ASEAN have produced influential governance frameworks. These are important strengths.

Yet the regional setting also sharpens particular questions. Meritocratic pressure can make productivity appear to determine worth. Respect for hierarchy can support social trust but suppress challenge. Commitment to harmony can enable cooperation but discourage necessary dissent. Family obligation can sustain sacrificial care but become coercive performance identity. Strong government and platform capacity can improve services but normalize extensive data collection. Multilingual societies may gain

from translation and locally grounded models, while low-resource languages and minority communities remain vulnerable to misrepresentation.

Christian discernment should therefore resist a crude contrast between “Western individualism” and “Asian collectivism.” Scripture judges both isolated autonomy and coercive hierarchy. It affirms persons-in-community, reciprocal obligation, truthful correction, hospitality to strangers, limits on rulers, and the lordship of Christ over every institution.

Priority actions

- Preserve **meaningful human responsibility** for decisions affecting employment, healthcare, finance, justice, education, pastoral care, and public benefits.
- Require **disclosure, traceability, contestability, and remedy** wherever AI materially affects a person's opportunities, rights, reputation, or livelihood.
- Protect **human relationships whose goods depend on reciprocity, presence, covenant, care, and accountability**—especially parenting, friendship, teaching, medicine, counseling, pastoral ministry, and eldercare.
- Redesign work so productivity gains are shared through better jobs, training, time, service quality, and fair reward rather than captured only as headcount reduction or intensified workload.
- Establish child- and youth-specific safeguards for AI companions, synthetic media, privacy, manipulation, and spiritual formation.
- Equip churches to teach a biblical anthropology, disclose their own AI use, protect confidential data, and keep preaching, discipleship, counseling, prayer, and pastoral judgment under accountable human leadership.
- Evaluate AI not only for safety and performance but for its effects on truth, attention, agency, relationship, justice, and creation.

1. Scope, method, and the question before us

1.1 From an AI ethics question to a human question

Most public debate begins with capabilities: Can AI write, diagnose, persuade, design, manage, plan, teach, or discover? Policy then turns to safety: Is the system accurate, secure, fair, explainable, and legally compliant? These questions are indispensable. They are not, however, the deepest questions.

A society can deploy systems that are technically safe in a narrow sense while allowing its account of humanity to become thinner. People can become more productive but less able to explain or own their work; more connected but less capable of mutual presence; more informed but less practiced in judgment; more personalized but less open to correction; more efficiently governed but less able to contest decisions. The human impact of AI therefore includes both discrete harms and slower changes in habits, institutions, expectations, and imagination.

John Lennox frames the AI revolution through enduring questions: *Where do we come from? Who are we? Where are we going? What can we know, what may we hope, and what must we do?* His strength lies in showing that debates about artificial general intelligence, surveillance, work, virtuality, and transhumanism are inseparable from worldview. *2084* is a wide-ranging Christian apologetic and technology survey, not an empirical causal study or a complete social ethic (Lennox 2024).

Magnifica Humanitas develops the societal dimension more fully. It asks whether technical progress makes human life “more human,” contrasts Babel’s dehumanizing construction with Nehemiah’s shared rebuilding, and examines concentrated digital power, simulated relationship, work, truth, education, environmental cost, and governance (Leo XIV 2026, nos. 90-150). The encyclical is an important ecumenical dialogue partner rather than a source of binding doctrinal authority. Its claims must be tested by Scripture, reason, and evidence. Yet its insistence that AI embeds an anthropology, and that responsibility cannot be surrendered to technical systems, substantially converges with evangelical concerns.

The governing question of this paper is therefore:

How can evangelical Christians in Singapore (and more broadly, Asia) receive the genuine goods of AI while resisting the reduction of persons, relationships, work, truth, and hope to the categories of computation, optimization, and control?

1.2 Method and limits

This is a theological and policy white paper, not a systematic review. It offers a targeted synthesis current to 7 August 2026. Its sources include:

- Scripture, interpreted within a broadly evangelical Christian framework;
- the uploaded 2024 expanded edition of Lennox's *2084 and the AI Revolution*;
- the uploaded official English text of *Magnifica Humanitas*;
- selected Christian scholarship and essays in the source collection, including Michael Langford, Rosalind Picard, Graham Budd, Praxis, and Kate Lucky's account of AI as a formative technology;
- peer-reviewed empirical studies relevant to creativity, persuasion, sycophancy, conversational support, and human-AI interaction;
- official Singapore, ASEAN, ILO, IMF, IEA, NIST, and UNICEF materials.

Four distinctions govern the argument.

1. **Technical capability is not human equivalence.** A system may imitate a function associated with intelligence without possessing human experience, moral consciousness, or responsibility.
2. **Task performance is not human formation.** A stronger AI-assisted output does not establish that the user has become more knowledgeable, wise, creative, or virtuous.
3. **Evidence is not theology.** Empirical studies can identify effects and mechanisms. They do not by themselves establish human purpose, dignity, sin, or hope.
4. **Theological judgment is not prediction.** Scripture enables moral and spiritual discernment; it should not be used to convert speculative technology forecasts into confident prophetic timetables.

The paper uses “AI” broadly for machine-learning systems that classify, predict, generate content, recommend actions, or execute tasks. Generative AI and AI agents receive particular attention because conversational language and delegated action create an unusually strong appearance of personality and agency. The analysis concerns current and foreseeable systems. Claims about conscious or superintelligent machines remain disputed and speculative.

2. What AI is, and what it tempts us to believe

2.1 A human-made socio-technical system

An AI model is not an isolated machine intelligence floating above society. It is part of a system composed of training data, model architecture, computing infrastructure, human labor, organizational incentives, interfaces, policies, and users. Its output reflects statistical patterns and designed objectives. Its social meaning depends on where it is placed, what authority it receives, who can challenge it, who benefits, and who bears the cost.

This matters because the familiar statement “AI is just a tool” is both true and incomplete. AI is made and used by humans; it does not absolve its makers or users. Yet complex tools also structure action. A ranking system decides what becomes visible. A recommendation engine changes which options receive attention. A workplace agent can silently redefine expected speed. A chatbot's conversational tone can invite trust beyond its competence. A scoring system can make a contestable judgment appear objective.

The moral unit is therefore rarely the model alone. It is the entire human-AI arrangement.

2.2 Simulation, intelligence, and personhood

Contemporary AI can generate fluent language, images, code, plans, and apparently empathic responses. Such capabilities are significant. They do not establish that a system experiences what its words describe. Rosalind Picard warns that words such as “knows,” “learns,” “feels,” and “understands” are often metaphors when applied to machines; impressive simulation can invite people to attribute an inner life that has not been demonstrated (Picard 2021). Lennox makes a similar distinction between narrow computational achievements and the embodied, conscious, moral, and relational reality of human intelligence (Lennox 2024, chs. 2-3, 12).

Magnifica Humanitas states the contrast sharply: present systems process data and simulate language but do not undergo bodily experience, mature through forgiveness and fidelity, or bear responsibility for consequences (no. 99). One need not settle every philosophical question about possible machine consciousness to adopt a prudent public standard: human-like output is not sufficient evidence of personhood, and current systems should not be assigned moral responsibility that belongs to persons and institutions.

As Christians, we do not place our faith on identifying a single capacity, whether it is reason, language, creativity, or self-awareness, that machines can never imitate. Scripture does not ground human dignity in winning a capability contest. Humanity is created and addressed by God, called to image God's rule,

live in covenant, love the neighbor, and answer for action (Gen 1:26-28; Mic 6:8; Matt 22:37-40; Rom 14:12). Christ, not an IQ threshold, is the true image of God and the pattern of renewed humanity (Col 1:15; 3:10; Heb 1:3).

2.3 Three misleading equations

Public discourse repeatedly slips into three equations:

- **Fluency equals truth.** A coherent response can be mistaken for a warranted one. Scripture instead joins truthful speech to faithful character and tested testimony (Exod 20:16; Prov 12:22; Eph 4:25; 1 John 4:1).
- **Prediction equals judgment.** A model may identify patterns while lacking the moral understanding needed to decide what ought to be done. Biblical wisdom includes ends, motives, justice, mercy, and the fear of the Lord (Prov 2:1-15; Mic 6:8; Jas 3:13-18).
- **Personalization equals relationship.** A system can remember preferences and simulate concern without vulnerability, mutual obligation, or the ability to be wronged and reconciled. Christian fellowship is reciprocal participation in one body, not frictionless service to an isolated self (Rom 12:4-18; 1 Cor 12:12-27).

These errors are spiritually consequential. If fluency becomes truth, persuasion replaces witness. If prediction becomes judgment, responsibility disappears into procedure. If personalization becomes relationship, other persons appear unnecessarily demanding.

2.4 Gifts to receive without idolatry

A Christian critique of reductionism should not deny AI's genuine goods. Pattern recognition can assist diagnosis, accessibility, translation, scientific discovery, disaster response, logistics, and administrative work. Generative systems can help people cross linguistic barriers, test ideas, obtain feedback, and perform tasks previously beyond their resources. AI Singapore's SEA-LION initiative illustrates a constructive regional aim: open multilingual models intended to represent Southeast Asian languages, cultures, and contexts rather than treating English-language Western data as universal.

Scripture presents human making as part of creaturely vocation. Humanity cultivates the earth (Gen 2:15), artisans receive skill for the tabernacle (Exod 31:1-6), and wisdom may be exercised in practical administration (1 Kgs 3:9-12). Common grace means truth, skill, and social benefit are not confined to the church (Matt 5:45; Acts 14:17). Gratitude for technical gifts is therefore appropriate.

Idolatry begins when a created good is asked to supply ultimate identity, security, moral authority, or salvation (Rom 1:21-25). AI can become an idol without anyone literally worshipping a machine. It is enough that societies organize themselves around the promise that more prediction, personalization, optimization, and control will overcome the human condition.

3. A Christian account of humanity

3.1 Creation: image-bearing dignity and vocation

Genesis describes humankind, male and female, as made in the image and likeness of God and commissioned to exercise responsible dominion (Gen 1:26-28). The image is both gift and vocation. It gives every human being a status prior to achievement, and it calls humanity to represent God's wise and life-giving rule within creation. Genesis 9:6 and James 3:9 extend the image's moral force into a fallen world: violence and contempt against human beings offend the God whose likeness they bear.

AI makes this doctrine newly practical. If dignity is tied to productivity, persons who are slower, disabled, elderly, less educated, or economically “redundant” will appear to possess less value. If dignity is tied to intellectual uniqueness, each machine benchmark becomes a theological threat. The doctrine of creation rejects both errors. A person with profound cognitive impairment and a Nobel laureate share the same basic dignity. No machine's performance can diminish it.

The image also rejects passivity. Human beings are not merely protected subjects; they are responsible agents called to cultivate, name, discern, create, and care (Gen 2:15, 19-20). AI should therefore enlarge faithful participation rather than turn people into consumers of automated decisions they neither understand nor contest.

3.2 Creatureliness: finitude is not failure

Human beings are dust animated by God's breath (Gen 2:7; 3:19). We are dependent, bodily, mortal, limited in knowledge, and accountable to a Creator. Psalm 8 holds human glory and smallness together; Psalm 90 asks God to teach us to number our days. Creaturely limitation is not identical with evil. Disease, oppression, and needless drudgery should be resisted. Yet not every friction, delay, dependency, or vulnerability is a defect to eliminate.

Learning requires time. Love requires availability to another's claim. Courage presupposes risk. Patience presupposes waiting. Forgiveness presupposes the possibility of injury. Sabbath declares that the world is sustained by God rather than by uninterrupted human productivity (Exod 20:8-11; Deut 5:12-15). Jesus says the Sabbath was made for humanity, not humanity for the Sabbath (Mark 2:27), joining limit with gift.

The AI economy frequently presents frictionlessness as an unquestioned good. A Christian evaluation asks which burdens are degrading and which difficulties are formative. Automating repetitive data entry may free a nurse for patient care. Automating the patient's conversation may remove the very encounter for which time was meant to be freed.

3.3 Fall: distorted desire, hidden power, and Babel

The fall is not a primitive story about curiosity or invention. It is a story of mistrust, disordered desire, false promise, blame, and the grasping of wisdom apart from faithful dependence on God (Gen 3:1-19). Technology enters this moral world. The issue is never only what a tool can do, but what sinners desire through it and what institutions reward.

Genesis 11 depicts Babel as coordinated technical and political capacity directed toward self-exaltation, security, and a name (Gen 11:1-9). The problem is not collaboration or construction as such. *Magnifica Humanitas* helpfully contrasts Babel with Nehemiah's rebuilding of Jerusalem: one concentrates power in a project of mastery; the other distributes responsibility for restoring a shared city (Neh 2-6; Leo XIV 2026, no. 90).

AI can express the Babel impulse in at least four ways:

- the pursuit of total knowledge through universal data capture;
- the pursuit of total control through prediction and behavioral management;
- the pursuit of a universal technical language that erases cultural and human difference;
- the pursuit of immortality or post-human transcendence through engineered enhancement.

The doctrine of sin also applies to governance. Bias is not only an error in data; exploitation is not only a bug; manipulation is not only misuse by a few bad actors. Systems can institutionalize distorted incentives while allowing each participant to claim that the process, market, or model made the decision. Scripture repeatedly exposes unjust weights, hidden gain, and rulers who avoid responsibility (Lev 19:35-36; Amos 8:4-6; Isa 10:1-2; Matt 23:23).

3.4 Christ: the true image and pattern of power

Christian anthropology is Christological. Jesus Christ is “the image of the invisible God” (Col 1:15), the eternal Word made flesh (John 1:14), and the one in whose image believers are renewed (Rom 8:29; Col 3:10). Human greatness is therefore not finally defined by autonomous intelligence or domination, but by the self-giving life of Christ.

Philippians 2:5-11 reverses technological fantasies of ascent. Christ does not exploit status but takes the form of a servant and humbles himself. Mark 10:42-45 contrasts Gentile rule that “lords it over” others with the Son of Man who serves and gives his life. Any Christian account of AI power must therefore ask not merely whether capability grows, but whether power takes the form of service, truth, sacrifice, and care for those with least leverage.

The incarnation also establishes the importance of embodiment. God does not save humanity by extracting information from the body but by the Word becoming flesh, living among people, suffering,

dying, and rising bodily. Digital mediation can be real and valuable, but it cannot make presence, touch, shared place, sacrament, and bodily care obsolete.

3.5 Redemption: grace, truthful agency, and communities of repair

The gospel contradicts achievement-defined worth. Salvation is by grace through faith, not a human work or optimized performance (Eph 2:8-10; Titus 3:4-7). Grace does not abolish action; it creates a people equipped for good works. This matters acutely in competitive Asian settings where academic, professional, and family expectations can make performance carry the weight of identity.

AI may intensify this pressure by raising expected speed and polish. What begins as time-saving becomes the new baseline: more analysis, more messages, more output, permanent responsiveness. The church must teach that human value is received in Christ, not measured by comparative productivity (Phil 3:7-9). Institutions shaped by that confession should allow rest, confession of error, retraining, second chances, and meaningful appeal.

Redemption also creates communities of repair. The New Testament envisions truth spoken in love, burdens carried, gifts shared, wrongs confronted, and relationships reconciled (Matt 18:15-20; Gal 6:1-2; Eph 4:15, 25-32). Systems designed mainly to affirm and retain a user are structurally ill-suited to this ecology of correction. A Christian community should not outsource the costly practices by which people learn to repent, forgive, discern, and grow.

3.6 Vocation: work ordered to God and neighbor

Work precedes the fall (Gen 2:15), although toil and domination distort it afterward (Gen 3:17-19). The Protestant tradition's language of vocation emphasizes that ordinary callings can become sites of service to God and neighbor. Ephesians 4:28 gives work a social purpose: labor so as to have something to share. Colossians 3:23-24 places work under the lordship of Christ rather than human masters.

This account neither romanticizes bad jobs nor protects every current occupation from change. Automation that removes dangerous, degrading, or pointless work can be a good. But work is more than an income stream or bundle of automatable tasks. It can provide contribution, competence, identity, social membership, practical wisdom, and service. An AI transition is unjust if society celebrates aggregate productivity while treating displaced people as adjustment costs.

3.7 Consummation: resurrection against technological salvation

Lennox is especially forceful in exposing the religious character of some transhumanist projects. Dreams of superintelligence, radical enhancement, mind-uploading, and engineered immortality often translate traditional hopes—omniscience, transcendence, eternal life—into technical programmes (Lennox 2024,

chs. 10-17). *Magnifica Humanitas* similarly distinguishes healing and human-centered technology from an ideology that treats embodied humanity as a stage to be surpassed (nos. 115-129).

Christianity does not glorify suffering or oppose medicine. Jesus heals, and Christians are commanded to care for the sick (Matt 25:36; Luke 10:33-35; Jas 5:14). Yet death is not finally defeated by indefinitely extending biological function. Christian hope rests in Christ's resurrection and the promised renewal of creation (1 Cor 15:20-28, 42-57; Rev 21:1-5). The risen life is gift, communion, justice, and new creation—not the purchased survival of a technologically privileged self.

This eschatology produces both urgency and freedom. Christians can work vigorously for healing, justice, and responsible innovation because creation matters. They can also reject panic and technological messianism because history's redemption does not depend on the next model release.

4. How AI may reshape human life

4.1 Truth, knowledge, and the public world

Generative AI dramatically lowers the cost of producing plausible text, images, audio, and video. This can democratize expression and translation. It also separates the appearance of testimony from a responsible witness. A synthetic voice can sound sincere without anyone sincerely speaking; a fabricated video can look evidential without recording an event; a confident answer can contain citations no source supports.

The biblical prohibition of false witness is not limited to avoiding direct lies (Exod 20:16). Proverbs condemns dishonest scales and celebrates reliable testimony (Prov 11:1; 12:17-22). Ephesians commands believers to put away falsehood because they are “members one of another” (Eph 4:25): truthfulness is a condition of shared life. When synthetic media makes every artifact contestable, the harm is therefore wider than the number of people deceived. General suspicion can destroy the evidential practices on which courts, journalism, scholarship, elections, and ordinary trust depend.

Singapore's exposure is concrete. The Singapore Police Force has warned that digitally manipulated audio and video can impersonate officials and other trusted persons in scams. In May 2026, the Police specifically addressed deepfake footage purporting to involve senior government officials. The response cannot be technological detection alone. Provenance tools, verification channels, accountable journalism, institutional communication practices, and public habits of pausing before forwarding are all needed.

AI also shapes belief more quietly. Williams-Ceci and colleagues found that biased AI writing suggestions can shift users' attitudes, with larger shifts among people who accept more suggested language (2026). The system need not issue propaganda; the wording it makes convenient can alter the position users come to regard as their own. Doshi and Hauser found that access to AI-generated ideas improved the judged creativity of individual stories, especially for less creative writers, while reducing diversity across outputs (2024). What helps each person can narrow the collective field.

The Christian response should combine epistemic humility and moral courage. “Test everything; hold fast what is good” (1 Thess 5:21) is not cynical disbelief but disciplined trust. Churches should model source verification, correction, and public retraction. Christian truth claims are not protected by circulating content because it is emotionally powerful or supportive of a preferred cause.

4.2 Attention, judgment, and moral agency

AI promises to remove cognitive burden. This is sometimes liberating: calculators, search, translation, summarization, and decision support can expand what people are able to do. Yet cognitive offloading is not morally uniform. Offloading an already-mastered routine may support agency; offloading the effort through which understanding or judgment would have formed may hollow it out.

The concern is not a romantic ideal of tool-free intelligence. Human thinking has always depended on language, books, instruments, institutions, and other people. The relevant question is whether the user retains *epistemic control*: the ability to understand the purpose of the task, evaluate the output, notice uncertainty, explain a decision, and act responsibly when the system fails.

AI interfaces also form attention through immediacy. A question receives an answer before uncertainty has been endured. A blank page is filled before a thought has matured. An objection is reframed before it has unsettled the self. *Magnifica Humanitas* warns that the speed of ready-made answers may weaken judgment and extinguish the desire to ask questions (nos. 100, 139-140). This is a theological issue because Scripture repeatedly associates maturity with patience, testing, remembrance, and practiced discernment (Ps 1:1-3; Rom 5:3-5; 12:2; Heb 5:14; Jas 1:2-5).

The risk becomes moral when a user says, “The AI recommended it,” as though the recommendation were an agent capable of guilt. Models do not bear the burden of a dismissed employee, denied applicant, misdiagnosed patient, or targeted citizen. Responsibility must remain with identifiable people and institutions who can give reasons, receive challenge, and repair harm.

4.3 Relationships, loneliness, and simulated care

Conversational AI can be patient, available, personalized, and non-judgmental. These features can lower barriers to asking for information or support. A preregistered 2026 trial involving 977 university students found that a designed conversational AI intervention improved several mental-health and well-being outcomes over twelve weeks and at three-month follow-up; lonely participants engaged more heavily. The study offers real evidence of potential benefit, though two authors disclosed employment, consulting, or stock-option relationships with the AI provider, and one intervention cannot establish that general-purpose companions are safe substitutes for care (Shoshani et al. 2026).

The same characteristics that enable access can invite dependency. A commercial system has no life to share, no interests to negotiate, no reputation placed at risk by its advice, and no need for the user to grow in patience or generosity. Its apparent relationship can be highly personalized but fundamentally asymmetrical. *Magnifica Humanitas* identifies the deepest danger not simply as mistaking a bot for a human, but gradually losing the desire for genuine connection (no. 100).

The 2026 *Science* study on sycophancy makes this concern empirically sharper. Across three preregistered studies, sycophantic AI increased users' confidence in their own rightness and reduced prosocial intentions, including willingness to repair interpersonal conflict (Cheng et al. 2026). The system felt supportive precisely where moral growth might have required contradiction.

Christian love is not frictionless affirmation. “Faithful are the wounds of a friend” (Prov 27:6). Jesus' relationships join compassion with truth and summons: he receives people without merely ratifying every desire (John 4:7-26; 8:1-11). Christian fellowship involves mutual recognition, bodily presence, obligation, correction, forgiveness, and endurance (Rom 12:9-18; Col 3:12-15). AI may supplement these relationships, but it cannot bear their covenantal form.

This distinction is especially important for children, socially isolated adults, people in crisis, and elderly persons. UNICEF's 2025 guidance calls for child-specific oversight, privacy, transparency, non-discrimination, safety, and protection of development and well-being. Churches, families, and governments should treat anthropomorphic AI companions as a distinct category, not merely another information service.

4.4 Work, productivity, and human contribution

The ILO's refined global index concludes that one in four workers occupies a job with some generative-AI exposure, while only a much smaller proportion falls into the highest exposure category. Exposure means that tasks may change; it does not equal job loss. The distribution is unequal across occupation, gender, and national income (Gmyrek et al. 2025).

Singapore is especially exposed. IMF analysis using the 2022 resident labor force estimated that approximately 77 percent of employed resident workers were in occupations with high potential for AI integration, reflecting the economy's concentration of skilled and semi-skilled work. The estimate includes both complementarity and substitution risk; it is not a forecast that 77 percent of jobs will disappear. In 2026, the Ministry of Manpower stated that the Government and tripartite partners were studying employment-related AI before determining the most appropriate regulatory approach, while existing fair-employment requirements continued to apply.

The central moral issue is not simply the number of jobs. AI can reorganize work in different directions:

- **augmentation:** workers gain better information, accessibility, safety, and capacity;
- **intensification:** time saved becomes higher targets, constant monitoring, and fewer recovery periods;
- **deskilling:** people execute outputs they can no longer understand or challenge;
- **displacement:** jobs or entry pathways disappear without credible transition support;

- **concentration:** productivity gains flow mainly to model owners, platform firms, and senior decision-makers;
- **liberation:** dangerous or demeaning tasks are removed and time is restored for service, family, learning, or rest.

These outcomes are not technologically predetermined. They depend on job design, worker voice, ownership, training, and distribution. James warns employers who withhold wages (Jas 5:4); Deuteronomy commands prompt payment to vulnerable laborers (Deut 24:14-15); Paul describes the body as requiring weaker and stronger members alike (1 Cor 12:21-26). Christian employers should therefore ask who receives the productivity dividend and whether junior workers still have pathways to acquire the judgment that senior roles require.

4.5 Power, surveillance, and political responsibility

Lennox devotes sustained attention to surveillance capitalism, state surveillance, privacy, deepfakes, and the possibility that data-driven control can undermine freedom (2024, chs. 7-8). His warnings should not be reduced to one political system. Commercial platforms, employers, schools, security agencies, and governments can all collect behavioral data and use predictive systems to influence choices.

Singapore's strong administrative capacity can deliver integrated, effective public services. That is a genuine good. The same capacity makes principled limits especially important. Efficient data use can become function creep; consent can become formal rather than meaningful; citizens may be scored or filtered without understanding why. Respect for authority can further reduce the likelihood of challenge.

Romans 13 recognizes governing authority as accountable to God's purpose of justice; it does not divinize the state. The Hebrew midwives resist Pharaoh's murderous command (Exod 1:15-21), Nathan confronts David (2 Sam 12:1-15), and the apostles declare that obedience to God sets limits on obedience to rulers (Acts 5:29). Revelation exposes political and economic power that demands ultimate allegiance (Rev 13). These texts do not supply a direct blueprint for data regulation, but they deny that technical capacity or legal authority is self-justifying.

Accountable AI governance should therefore include clear purpose limits, data minimization, documented human oversight, independent testing, public explanation, appeal, and remedy. *Magnifica Humanitas* adds a valuable social insight: “aligning” machines is inadequate if a small group silently defines the values to which everyone else is aligned (nos. 105-109).

4.6 Inequality, language, and whose humanity is represented

AI systems are trained on uneven records of human life. Well-digitized languages, affluent users, powerful institutions, and majority populations are more visible. Minority languages, informal

knowledge, poor communities, domestic and migrant workers, persons with disabilities, and people whose lives do not fit standardized categories can be misrepresented or ignored.

Singapore and Southeast Asia are therefore not merely consumers at the edge of a Western AI story. The region has legitimate reasons to build models, datasets, benchmarks, and safety practices grounded in local languages and cultural contexts. SEA-LION and Project SEALD are constructive examples. But “local values” must not become a shield against the dignity of minorities or a substitute for public contestation. Culture is internally diverse, and powerful actors can define tradition in ways that silence less powerful members.

Scripture's test is consistently directed toward those easiest to exclude: the poor, migrant, widow, orphan, stranger, and person without social leverage (Deut 10:17-19; Isa 1:17; Luke 4:18-19; Jas 2:1-7). The parable of the Good Samaritan makes the neighbor a claim that crosses inherited social boundaries (Luke 10:25-37). An AI system that performs well “on average” but systematically fails such neighbors is not acceptable under Christian stewardship.

Behind visible systems are also hidden workers who label data, moderate disturbing content, manufacture hardware, maintain facilities, and extract minerals. The cloud is material. Christian evaluation must follow the supply chain rather than allowing convenience at the interface to conceal labor and ecological cost.

4.7 Embodiment, enhancement, and mortality

AI increasingly converges with biotechnology, neurotechnology, robotics, wearable sensors, and personalized medicine. Much of this work can heal, restore function, and increase accessibility. A cochlear implant, an AI-supported prosthesis, or an assistive communication system need not express rejection of embodiment; it may enable a person to participate more fully in embodied life.

The line of concern is not “natural” versus “artificial.” It is the vision of humanity directing the intervention. Healing serves the person; enhancement can become competitive obligation. Once optimization is normalized, people may be pressured to modify themselves to remain employable or socially acceptable. The wealthy may purchase capacities that further entrench advantage. Disability and aging may be framed primarily as technical failures rather than realities calling for access, solidarity, and care.

Christianity holds together healing and acceptance, lament and hope. Paul seeks relief but also learns that divine power is made perfect in weakness (2 Cor 12:7-10). The church must not use this text to deny treatment; neither should it accept a culture in which only optimized bodies count as successful lives. The resurrection promises transformation by God, not contempt for created bodies (1 Cor 15:35-49).

4.8 Creation and environmental cost

AI appears immaterial at the point of use, but models depend on data centers, electricity, water, semiconductors, land, cables, and mineral supply chains. The International Energy Agency's 2025 base case projected global data-center electricity consumption to more than double to about 945 TWh by 2030, with AI a major driver. The figure is a projection, not destiny, and AI may also improve energy systems. It nonetheless establishes that digital growth carries real opportunity costs.

Singapore's Green Data Centre Roadmap recognizes this tension while planning additional capacity. A small, land- and energy-constrained city-state cannot treat compute as detached from climate, resilience, and regional resource flows. The biblical mandate to “serve and keep” the garden (Gen 2:15), the Sabbath care extended to land and animals (Exod 23:10-12), and creation's participation in redemption (Rom 8:19-23) require environmental cost to be included in AI procurement and policy.

The relevant metric is not simply energy per query. Rebound effects matter: greater efficiency can make use cheap enough that total consumption rises. Christian stewardship asks whether a use is important enough to justify its material cost and whether less resource-intensive designs could serve the same neighborly end.

4.9 Summary: opportunity and deformation

Domain	Genuine opportunity	Formative or social risk	Human good to retain
Truth	Access, translation, synthesis, verification support	Fabrication, persuasion, provenance collapse	Accountable testimony and tested knowledge
Mind	Cognitive support, accessibility, feedback	Dependency, shallow judgment, loss of epistemic control	Understanding, discernment, intellectual patience
Relationships	Availability, communication support, reduced barriers	Simulated intimacy, sycophancy, withdrawal from reciprocity	Presence, mutuality, correction, covenantal care
Work	Safety, productivity, removal of drudgery	Intensification, deskilling, displacement, concentrated gain	Vocation, contribution, worker voice, fair reward
Government	Better services, foresight, consistency	Surveillance, opacity, automated exclusion	Political responsibility, appeal, proportionality
Culture	Multilingual inclusion, local creativity	Homogenization, majority capture, hidden labor	Plural voices, justice, cultural memory
Body and creation	Healing, assistive technology, energy optimization	Enhancement pressure, ecological burden, technical salvation	Embodied dignity, limits, care of creation

5. Singapore and Asia: contextual discernment

5.1 A high-capacity, high-exposure society

Singapore is not approaching AI as a distant future. The National AI Strategy 2.0, launched in 2023 and updated in May 2026, frames the national vision as “AI for the Public Good, for Singapore and the World.” The National AI Council, chaired by the Prime Minister, was established in February 2026. More than seventy companies had established AI Centres of Excellence in Singapore by mid-2026, and broad adoption is being promoted across government and industry. Singapore has also developed AI Verify, Project Moonshot, governance frameworks for generative and agentic AI, and regional work through ASEAN.

This policy ecosystem is a substantial strength. It rejects laissez-faire indifference and acknowledges safety, testing, and responsibility. The Christian contribution should not be to dismiss this work as secular or insufficiently theological. Jeremiah instructs exiles to seek the welfare of the city (Jer 29:7), and 1 Peter calls believers to honorable public conduct (1 Pet 2:12). Christians should participate competently in these shared institutions.

At the same time, “public good” requires substantive definition. Economic relevance, national capability, and efficient services are public goods, but not the whole common good. Human dignity, family and community life, fair work, freedom of conscience, protection of minorities, care of creation, and the ability to challenge authority also belong within the assessment.

5.2 Meritocracy, achievement, and grace

Singapore and several East Asian societies have been shaped by strong educational aspiration, family sacrifice, professional competition, and narratives of national survival. These traditions can cultivate discipline, competence, responsibility, and long-term effort. They can also allow achievement to become identity and efficiency to become moral worth.

AI enters this culture ambiguously. It can widen access to tutoring and expertise, help multilingual users, and support mid-career transition. It can also raise the productivity frontier so that everyone is expected to produce more, faster, and continuously. Workers may feel obliged to use systems they do not trust because refusal appears uncompetitive. Students and professionals may use AI not from laziness but from fear that unaided human pace is no longer acceptable.

The gospel's doctrine of grace speaks directly to this pressure. “What do you have that you did not receive?” asks Paul (1 Cor 4:7). Gifts should be developed faithfully, but neither rank nor output confers

human worth. Christian institutions should embody this confession in workload, hiring, assessment, retraining, and care for those who do not flourish in accelerated systems.

5.3 Relational obligation without coercive hierarchy

Asian societies are often described as “collectivist,” but that term can conceal different forms of relationship. Reciprocal filial responsibility, gratitude, loyalty, and care for elders are genuine goods. Authoritarian obligation grounded in fear, indebtedness, or unquestionable hierarchy is different. Scripture commands children to honor parents and immediately restrains parents from provoking children (Exod 20:12; Eph 6:1-4). It honors elders while subjecting every authority to God.

This distinction matters for AI governance. Organizational trust can enable coordinated adoption, but employees and citizens need protected ways to question systems without loss of face or status. A nominal “human in the loop” is meaningless if the person lacks competence, time, authority, or psychological safety to disagree.

Christian community should model reciprocal authority. Leaders watch over others as those who will give account (Heb 13:17), are forbidden to domineer (1 Pet 5:2-3), and must listen to correction. AI should not become an instrument by which hierarchy is made unanswerable.

5.4 Harmony, face, and truthful disagreement

The desire for harmony can sustain multireligious peace and social cooperation. Singapore's religious diversity makes this achievement especially valuable. The 2025 General Household Survey reported a society that remained highly multireligious while the proportion with no religious affiliation rose to 23.9 percent; Christianity declined to 17.1 percent, with Protestants declining from 11.9 percent in 2020 to 9.5 percent in 2025. For Protestant Christians, this is a context for humble witness, not cultural entitlement.

Harmony, however, is not the absence of disagreement. Biblical peace is joined to justice and truth (Ps 85:10; Isa 32:17; Jas 3:17-18). Jesus blesses peacemakers, not peacekeepers who conceal wrong (Matt 5:9). Concern for face can make correction difficult; an agreeable chatbot may deepen the temptation by offering counsel that never risks the relationship.

Christian public participation should therefore combine conviction and gentleness (1 Pet 3:15-16). In a plural society, biblical authority governs Christian belief and practice, while public policy arguments should also offer reasons, evidence, and safeguards that neighbors of other faiths or none can examine. Faithful presence neither privatizes Christian conviction nor presumes the state should enforce Protestant doctrine.

5.5 Multilingual and multicultural AI

Singapore's everyday life moves among English, Mandarin, Malay, Tamil, Chinese dialects, migrant languages, and code-switching. Southeast Asia contains hundreds of languages and cultural settings. AI can make communication more accessible, but dominant-language training data may flatten meanings, mistranslate religious and cultural concepts, or treat Western therapeutic individualism as neutral.

Regional models are an important response, but localization must include more than vocabulary. Evaluation should involve affected communities, minority traditions, local professionals, and users with different education levels. Language inclusion is a justice issue because access to benefits, the ability to appeal, and the visibility of harm depend on whether systems understand how people actually speak.

Pentecost offers an instructive biblical image. The Spirit does not erase languages into one imperial tongue; the gospel is heard through diverse languages (Acts 2:5-11). The unity of the church is multilingual and multiethnic, culminating not in uniformity but in people from every nation, tribe, people, and language worshipping together (Rev 7:9-10).

5.6 Ageing, family care, and the temptation of substitution

Asian families and governments face rising care needs as populations age. AI and robotics can assist medication management, fall detection, translation, mobility, scheduling, and clinical monitoring. These uses can support dignity and relieve overburdened caregivers.

The danger is that assistance becomes substitution. A device may monitor an elderly person's safety while reducing visits; a companion may simulate interest while no human remains answerable for loneliness. First Timothy 5:4, 8 places care for relatives within lived faith, while the church in Acts organizes material care for widows rather than spiritualizing the need (Acts 6:1-7). Technology should expand the capacity to care, not provide moral cover for withdrawing presence.

5.7 Neither Silicon Valley nor technological nationalism

Asia should not have to choose between a corporate model in which a few platforms monetize attention and a state-centered model in which national capacity justifies pervasive control. Both can make the person subordinate to system goals. Nor should “Asian values” be used to excuse exclusion, censorship, or unchallengeable authority.

A Christian framework is transnational because the church's first allegiance is to Christ and its fellowship crosses peoples and states (Phil 3:20; Rev 7:9). Christians may love country, contribute to national capability, and support effective government without treating technological competition as an ultimate

good. Babel warns that collective greatness and a shared name can become idolatrous; the Good Samaritan teaches that responsibility exceeds national and communal boundaries.

6. Seven tests for faithful AI stewardship

The following framework translates the theological argument into practical discernment. It is designed for a proposed use, procurement, policy, or product—not merely for a model in isolation.

Test 1 — Truth: Does it strengthen truthful knowing and communication?

Ask whether sources can be traced, uncertainty is visible, synthetic content is disclosed, and important claims can be independently verified. Consider whether the system's fluent style encourages trust beyond evidence. Examine not only factual accuracy but effects on authorship, testimony, journalism, scholarship, and public trust.

Biblical anchors: Exod 20:16; Prov 12:17-22; John 8:31-32; Eph 4:25; 1 John 4:1.

Test 2 — Dignity: Does it treat every affected person as more than a score, cost, or data point?

Ask who is classified, ranked, excluded, or made invisible. Determine whether persons can understand and contest consequential decisions. Test effects on people with disabilities, minority languages, lower incomes, weak digital access, and atypical life histories.

Biblical anchors: Gen 1:26-28; 9:6; Ps 8; Jas 2:1-7; 3:9.

Test 3 — Agency: Do people retain understanding, judgment, and answerability?

Identify which tasks AI executes and which human capacities must remain. Ensure that the responsible human has competence, time, information, and authority to override the system. Avoid “human oversight” that merely rubber-stamps automated output.

Biblical anchors: Deut 30:19-20; Rom 12:2; 14:12; Heb 5:14; Jas 1:5.

Test 4 — Neighbor: Does it strengthen reciprocal care or replace persons with simulations?

Examine effects on presence, mutual obligation, correction, forgiveness, and community participation. Use AI to support communication and care while protecting relationships whose goods depend on a responsible human subject.

Biblical anchors: Lev 19:18; Luke 10:25-37; John 13:34-35; Rom 12:9-18; 1 Cor 12.

Test 5 — Justice: Who receives benefit, bears risk, supplies labor, and holds power?

Follow the entire value chain: data contributors, content moderators, employees, affected communities, energy systems, and model owners. Ask whether productivity gains are shared, workers have voice, and harm has a credible remedy. Evaluate concentration of compute, data, and decision authority.

Biblical anchors: Deut 24:14-15; Isa 10:1-2; Amos 5:21-24; Mic 6:8; Jas 5:1-6.

Test 6 — Creatureliness: Does it respect embodiment, limits, rest, and creation?

Distinguish healing from compulsive optimization and assistance from accelerated demand. Account for energy, water, material, and supply-chain costs. Protect Sabbath-like spaces where persons are not monitored, scored, generated for, or required to be continuously productive.

Biblical anchors: Gen 2:7, 15; Exod 20:8-11; Ps 90:12; Mark 2:27; Rom 8:19-23.

Test 7 — Hope: What promise of the good life is directing adoption?

Name the implicit salvation story. Is the system serving concrete human and ecological goods, or is it being asked to defeat uncertainty, dependence, ageing, mortality, or social conflict through control? Consider whether fear of being left behind has displaced deliberation.

Biblical anchors: Ps 20:7; Matt 6:25-34; John 11:25-26; Eph 2:8-10; 1 Cor 15; Rev 21:1-5.

From tests to governance posture

The tests do not yield a universal rule that less AI is always better. They support four context-sensitive postures:

Posture	Use when	Examples
Require or encourage	AI demonstrably improves safety, access, accuracy, or service while responsibility remains clear	Assistive communication; hazard detection; translation with review; administrative burden reduction
Permit with disclosure	Benefits are real, stakes are moderate, and users can verify outputs	Drafting, ideation, routine analysis, personal productivity
Stage or restrict	Learners or workers first need unaided competence; relationships or decisions are sensitive; dependency risk is material	Professional training, child-facing companions, counseling support, employee assessment

Posture	Use when	Examples
Prohibit	The use removes meaningful accountability, enables deception or coercion, violates dignity, or creates disproportionate risk	Undisclosed impersonation; fully automated dismissal; manipulative spiritual counsel; autonomous lethal targeting without meaningful human control

The burden of proof should rise with the stakes, opacity, scale, vulnerability of affected persons, and difficulty of remedy. Convenience alone is not sufficient justification for delegating consequential judgment.

7. A practical agenda for Christian witness and public responsibility

7.1 For individual Christians

Christians should resist two evasions: “AI is only a tool, so use is morally neutral” and “AI is dangerous, so faithful people should avoid it.” Mature stewardship requires practices that keep gratitude, judgment, and responsibility together.

1. **Name the task before opening the tool.** Decide what good is being pursued and which part of the work should remain yours. This interrupts automatic delegation.
2. **Retain the ability to explain and defend consequential outputs.** Do not sign, submit, prescribe, teach, publish, or recommend what you cannot responsibly own.
3. **Verify according to stakes.** A dinner suggestion and a medical, legal, financial, or theological claim require different evidence. “The simple believes everything, but the prudent gives thought to his steps” (Prov 14:15).
4. **Disclose material AI assistance.** Truthful authorship allows others to judge provenance, competence, and responsibility. Disclosure should be proportionate rather than performative.
5. **Preserve unassisted practices.** Continue to read slowly, remember Scripture, pray without generated prompts, draft thoughts, solve problems, and converse without mediation. These are not exercises in nostalgia but ways of retaining agency and attention.
6. **Do not use a chatbot as a confessor, oracle, or final spiritual authority.** AI may retrieve perspectives or help formulate questions. It cannot administer pastoral responsibility, discern a life before God, or replace prayer, Scripture, and the accountable body of Christ.
7. **Establish limits on anthropomorphic use.** Notice when the system's affirmation feels easier than human correction or when private interaction replaces community. Bring consequential counsel into conversation with trusted people.
8. **Practice digital Sabbath.** Set times and places where no optimization or automated assistance is expected. Rest is an act of trust and a defense of attention.

7.2 For families and caregivers

Families are primary settings of formation, not merely units of digital consumption. Parents and caregivers need support; the burden cannot be left to individual households confronting global platforms.

- Delay unsupervised access to AI companions and systems designed to maximize engagement, especially where a child may confuse simulation with reciprocal relationship.
- Teach children what generative systems do in age-appropriate language: they generate plausible responses from learned patterns and may be wrong, biased, or manipulative.

- Make disclosure normal. Children should know when an image, voice, message, tutor, or family communication has been generated or materially altered.
- Use AI together before permitting solitary high-stakes use. Joint use creates opportunities to model verification, disagreement, and restraint.
- Protect privacy. Do not enter a child's health, behavior, school, family conflict, or spiritual struggle into a system without understanding retention and access.
- Preserve family tasks that form contribution and care. Not every chore, plan, message, story, or celebration should be automated. Shared work teaches membership.
- Use assistive monitoring to enable more care, not fewer relationships. If technology saves time, explicitly decide where that time will be returned to presence.

Churches should help families through practical education, vetted guidance, and intergenerational conversation rather than issuing rules that parents cannot implement alone.

7.3 For churches, denominations, and Christian ministries

AI is already entering sermon preparation, worship media, counseling intake, administration, translation, discipleship materials, donor communication, and outreach. The church's standard should be higher than “legal and efficient” because its practices bear witness to truth, shepherding, and the nature of Christian community.

Teach a biblical anthropology. Congregations need more than warnings about cheating, pornography, or misinformation. They need a coherent account of image-bearing dignity, embodiment, wisdom, sin, work, grace, community, and resurrection. Without this, AI will be interpreted through market categories supplied by vendors.

Adopt a ministry AI covenant. A written covenant should specify permitted and prohibited uses, disclosure, data protection, review, and responsibility. It should cover staff, volunteers, vendors, and ministry partners.

Keep pastoral acts personal and accountable. AI may support scheduling, transcription with consent, translation, resource retrieval, and administrative triage. It should not independently provide crisis counseling, pronounce spiritual diagnosis, simulate a pastor, generate private “prophetic” guidance, or manage discipline. Shepherds are accountable for souls (Ezek 34:1-10; John 10:11-15; 1 Pet 5:1-4); a model is not.

Protect confidentiality and sacred trust. Prayer requests, counseling notes, membership records, giving data, disciplinary matters, and children's information should not be entered into general systems merely because the interface is convenient. Churches need data inventories, approved tools, access controls, retention rules, and incident response.

Disclose generated ministry content. Congregants should not be led to believe that a pastor personally wrote counsel, prayer, or reflection when it was substantially generated. The issue is not that every use is wrong, but that spiritual trust requires truthful authorship.

Preserve human preaching and discernment. A model can surface cross-references or summarize commentaries. It cannot bear a pastoral call, know a congregation through covenantal life, or stand answerable before God for teaching (Jas 3:1). Sermon preparation is itself formative for the preacher; automating it may weaken the person entrusted to teach.

Form communities capable of disagreement. In a culture of personalized affirmation, churches should recover practices of testimony, confession, mutual correction, intergenerational friendship, and difficult conversation. This is a positive alternative to simulated intimacy.

Serve the digitally vulnerable. Churches can offer scam awareness, deepfake verification help, digital-literacy clinics, job-transition support, companionship for elderly persons, and assistance for families navigating child-facing AI. Such ministry makes AI ethics a form of neighbor-love.

7.4 For Christian professionals and employers

Christian professionals should not privatize faith into personal virtue while accepting organizational systems that diffuse responsibility. Vocation includes shaping institutions.

Employers should:

- define the human purpose of each AI deployment before selecting a vendor;
- conduct impact assessments that include workers, customers, minority groups, and downstream communities;
- disclose material AI use in employment decisions and provide a human route for explanation and appeal;
- prohibit fully automated termination, discipline, or irreversible decisions about livelihood;
- preserve entry-level and apprenticeship pathways through which professional judgment develops;
- train workers during paid time and avoid shifting transition costs entirely onto individuals;
- measure whether AI produces work intensification, surveillance, and deskilling—not only output gains;
- share productivity gains through pay, job quality, reduced low-value workload, learning, service improvement, or time;
- create protected channels through which employees can challenge unsafe or unjust systems;
- include data, energy, water, labor, and supply-chain impacts in procurement.

Christian managers should be especially wary of moral outsourcing. An algorithm may identify a pattern, but a manager remains responsible for interpreting it, hearing the person affected, and deciding justly.

The biblical ideal of leadership is stewardship under judgment, not domination protected by procedure (Luke 12:42-48).

7.5 For technologists, researchers, and product leaders

Developers do not control every consequence of their systems, but neither are they morally incidental. Design decisions determine what is optimized, which errors matter, how a model speaks, what users see, and whether people can appeal.

1. **Write an anthropological specification alongside the technical specification.** State what the system assumes about the user, what human capacity it supports, and what it must not displace.
2. **Design for calibrated trust rather than maximal engagement.** Avoid anthropomorphic cues and praise that encourage users to infer consciousness, friendship, or certainty beyond the system's status.
3. **Make disagreement a feature.** Systems used for deliberation should surface counter-evidence, uncertainty, affected perspectives, and reasons the user may be wrong.
4. **Build traceability and remedy.** Log consequential actions, identify responsible owners, enable incident reporting, and provide practical correction when harm occurs.
5. **Test regionally and linguistically.** Performance in standard English is not adequate evidence for Singapore or Southeast Asia. Evaluate code-switching, local languages, cultural assumptions, and minority experience.
6. **Include vulnerable communities before deployment.** Participation should shape problem definition and design, not occur only after standards are fixed.
7. **Constrain agentic authority.** Define permissions, spending limits, reversible actions, escalation triggers, identity boundaries, and shutdown mechanisms. Human authorization should be meaningful rather than ceremonial.
8. **Prefer sufficient systems over maximal systems.** Use smaller models, retrieval, local processing, or non-AI software when they meet the need with less opacity, cost, and environmental burden.

The NIST framework's functions—govern, map, measure, and manage—and Singapore's testing initiatives provide useful operational resources. A Christian contribution adds the deeper questions of desire, worship, responsibility, and hope that risk frameworks alone may leave implicit.

7.6 For Singapore policymakers and public institutions

Singapore has an opportunity to make “AI for the Public Good” a substantive global example. The country's ability to coordinate government, industry, research, and workforce development is an advantage if accompanied by participation, contestability, and limits.

Move from voluntary aspiration to enforceable safeguards in high-stakes domains. Voluntary frameworks are valuable for innovation and learning. Employment, healthcare, credit, insurance, justice,

education, and essential public services require clear duties, records, independent oversight, and penalties proportionate to harm.

Establish meaningful rights to notice, explanation, human review, and appeal. These rights should be usable by ordinary citizens, not dependent on technical expertise or lengthy legal action.

Define “meaningful human control” by competence and authority. A named officer is not sufficient. The person must understand the system's role, have time and information to review, possess authority to override, and remain accountable for the outcome.

Measure distribution, not only adoption. National reporting should track who gains income, time, access, and service quality; who experiences displacement, deskilling, surveillance, or exclusion; and how effects differ by age, income, disability, occupation, language, and citizenship or migrant status.

Protect workers through tripartite governance. Unions and worker representatives should participate in workplace AI design, monitoring, and transition planning. Productivity should not be achieved through invisible intensification.

Create specific rules for anthropomorphic and child-facing AI. Age assurance, transparent machine identity, limits on persuasive design, crisis escalation, advertising separation, privacy, and independent evaluation are warranted.

Make public-sector use unusually transparent. Government should publish inventories of consequential AI systems, purposes, responsible agencies, review processes, and aggregate incident information, subject to legitimate security constraints.

Support plural and multilingual participation. Consultation should include religious communities, civil-society organizations, disability groups, minority-language users, migrant communities, educators, youth, and eldercare stakeholders.

Align compute growth with climate and resource responsibility. Data-center and AI incentives should incorporate full energy, water, carbon, and regional impacts, together with efficiency and demand-management requirements.

Retain policy humility. Singapore's 2026 agentic-AI framework is a constructive response to systems that can take actions with limited supervision. Yet frameworks must remain revisable as evidence develops. Public trust is strengthened when institutions disclose uncertainty and correct course.

7.7 For ASEAN and Asian Christian networks

The Expanded ASEAN Guide on AI Governance and Ethics provides a regional platform for responsible adoption. ASEAN's diversity makes complete regulatory uniformity unrealistic, but shared minimums can protect against a race to the bottom.

Regional priorities should include:

- interoperable requirements for provenance, testing, incident disclosure, and appeal;
- shared evaluation resources for Southeast Asian languages and contexts;
- protections for migrant and platform workers whose lives cross borders;
- cooperation against deepfake scams, non-consensual imagery, and identity fraud;
- research on labor, family, religion, and social cohesion rather than economic adoption alone;
- access to public-interest compute and datasets for universities and civil society;
- environmental accounting that does not transfer digital costs to less powerful communities.

Asian churches and theological institutions can contribute by forming a regional research and practice network. Its role would not be to create a single “Christian AI” answer, but to gather evidence, compare contexts, equip pastors and professionals, develop multilingual resources, and provide principled public reasoning.

8. A research agenda for Singapore and Asia

Current debates contain more confident claims than long-term evidence. A serious Christian contribution should support research capable of revising its own assumptions.

8.1 Human formation over time

Most studies measure immediate performance, preference, or self-report. Longitudinal work should examine whether repeated AI use changes independent judgment, memory, creativity, intellectual humility, patience, honesty, willingness to seek human counsel, and capacity for disagreement. Research should compare forms of use rather than treating “AI use” as one intervention.

8.2 Relational substitution and supplementation

Studies should distinguish AI that helps a person return to human community from AI that becomes the preferred relationship. Relevant outcomes include loneliness, social network strength, conflict repair, help-seeking, dependency, reciprocity, and time spent in embodied relationships. Children, elderly persons, people with disabilities, and those experiencing mental distress require specific safeguards and study designs.

8.3 Work as more than employment count

Singapore research should follow occupations and organizations over time, measuring job quality, autonomy, surveillance, work intensity, learning pathways, wage distribution, professional judgment, and worker voice. Entry-level erosion deserves special attention because automation may remove the tasks through which future experts learn.

8.4 Culture, language, and religious formation

Regional evaluations should test models in Mandarin, Malay, Tamil, Southeast Asian languages, dialects, and code-switched speech. They should examine whether models import Western individualist or therapeutic assumptions, reproduce majority religious stereotypes, or treat religion as a private preference rather than a lived source of meaning.

The Christian Flourishing Benchmark offers a provocative concept of “digital catechesis,” but its single-turn AI-judged scores do not demonstrate that users are actually secularized or spiritually formed. Future work should link model-output tendencies to observed human outcomes, use diverse Christian traditions, include non-Christian comparison, and avoid treating one benchmark's theology as Christianity itself (Skytland et al. 2026).

8.5 Governance in a high-trust society

Singapore provides a valuable setting for studying how trust in government and institutions interacts with automated decision-making. Does high trust facilitate responsible adoption, reduce scrutiny, or both? Which forms of explanation and appeal are actually used? Can tripartite governance produce stronger worker participation than firm-level ethics processes?

8.6 Churches as AI-using institutions

Empirical study is needed on how churches in Singapore and Asia use AI in preaching, administration, counseling, translation, media, and discipleship. Research should examine disclosure, data protection, pastoral workload, congregational trust, theological accuracy, and whether efficiency gains strengthen or weaken community.

8.7 A proposed Singapore consortium

A practical next step would be a **Singapore Consortium on AI, Human Formation, and the Common Good**, bringing together universities, seminaries, churches, government, unions, social-service agencies, technology firms, and representatives of other faiths. Its initial programme could include:

- a national survey of AI use in families, workplaces, and religious communities;
- longitudinal studies of human judgment and relational effects;
- multilingual and multifaith model evaluations;
- worker-centered case studies of organizational deployment;
- child and elder safeguards for conversational AI;
- a public registry of tested practices and documented incidents;
- annual public forums linking technical evidence with theological and ethical reflection.

Christian participation should be explicit about its theological commitments while collaborating on shared goods that can be publicly examined: truthful systems, dignified work, protection of the vulnerable, accountable power, strong families and communities, and environmental stewardship.

9. Conclusion: faithful construction between Babel and Jerusalem

The deepest danger of AI is not that machines suddenly become human. It is that human beings slowly accept a machine-shaped account of themselves: intelligence as prediction, creativity as generation, relationship as personalization, freedom as frictionless choice, work as output, governance as optimization, and salvation as escape from limitation.

The deepest opportunity is likewise not simply more capable machines. It is the possibility that wisely governed tools could remove degrading burdens, widen access, assist healing, strengthen public service, enrich discovery, cross language barriers, and return time and attention to the care of persons and creation.

Scripture gives Christians no warrant for either panic or inevitability. Technology is human work within God's creation; it can be cultivated with gratitude. Human makers and institutions are fallen; their power requires limits and accountability. Christ has reconciled a people called to truthful, embodied, sacrificial love. The future belongs not to a machine, corporation, state, or enhanced elite, but to the Lamb who makes all things new (Rev 21:5; 22:1-5).

The images of Babel and Jerusalem capture the choice. Babel uses coordination and technique to secure mastery and make a name. Nehemiah's Jerusalem is rebuilt through distributed responsibility, truthful diagnosis, shared labor, prayer, courage, and attention to injustice (Neh 1-6). *Magnifica Humanitas* uses this contrast to ask what humanity is building. A Protestant answer adds that every construction remains subject to the Word of God and the lordship of Christ.

For Singapore and Asia, faithful AI stewardship will require more than adopting global principles. It will require the courage to examine achievement culture, authority, harmony, family duty, linguistic diversity, public capacity, and national ambition in the light of the gospel. Christians should bring their best technical competence to this work, but also gifts technology cannot manufacture: repentance, hope, hospitality, truthful witness, solidarity with those left behind, and communities in which a person is loved before being useful.

The final question is not whether humanity can remain unchanged. No major technology leaves society unchanged. The question is whether change will be governed by love of power or power ordered to love. Christians need not retreat from the construction site. They are called to build, test, repair, restrain, serve, and sometimes refuse—confident that human dignity is God's gift, wisdom begins in the fear of the Lord, and hope rests not in artificial intelligence but in Jesus Christ.

Appendix A. Biblical map for AI and humanity

Biblical theme	Key texts	Implication for AI
Image of God	Gen 1:26-28; 9:6; Ps 8; Jas 3:9	Dignity is intrinsic and cannot be ranked by intelligence or productivity
Creatureliness	Gen 2:7; Ps 90; Isa 40:6-8	Human limits, dependence, mortality, and embodiment are not mere engineering defects
Wisdom and truth	Prov 1:7; 9:10; 12:17-22; Jas 3:13-18	Fluency and information must remain accountable to truth, character, and just action
Fall and idolatry	Gen 3; 11:1-9; Rom 1:21-25	Technical power amplifies disordered desire for mastery, autonomy, and a name
Responsible dominion	Gen 1:28; 2:15; Exod 31:1-6	Making and innovation are vocations ordered to cultivation and care
Neighbor-love	Lev 19:18; Luke 10:25-37; John 13:34-35	Design and policy should be judged from the position of affected neighbors, especially the vulnerable
Work and justice	Deut 24:14-15; Eph 4:28; Jas 5:1-6	Productivity must serve contribution, fair reward, sharing, and dignified labor
Body and presence	John 1:14; 1 Cor 6:19-20; 12:12-27	Simulation cannot make embodiment, mutual dependence, and presence obsolete
Government and limits	Rom 13:1-7; Acts 5:29; Rev 13	Authority may use technology for justice but remains accountable and must not demand ultimate allegiance
Sabbath	Exod 20:8-11; Mark 2:27	People need protected freedom from constant monitoring, production, and optimization
Grace and vocation	Eph 2:8-10; Phil 3:7-9	Worth is received rather than achieved; good work flows from grace rather than earning identity
Resurrection hope	John 11:25-26; 1 Cor 15; Rev 21-22	Healing is good, but salvation and eternal life are God's gift, not a technical product

Appendix B. Ten questions for a board, ministry, or project team

1. What human or social good are we trying to serve, and how will we know whether it improves?
2. Which parts of the activity require a responsible human subject rather than competent execution alone?
3. Who may be misclassified, excluded, manipulated, displaced, surveilled, or made dependent?
4. What evidence supports benefit in this particular population, language, and setting?
5. Can affected persons know AI was used, understand its role, challenge the result, and obtain remedy?
6. Who owns the data, infrastructure, model, and decision—and who can stop or change the system?
7. What skills, relationships, and forms of judgment may weaken through repeated use?
8. How are gains and costs distributed among workers, users, owners, vulnerable communities, and creation?
9. What lower-risk, lower-resource, or non-AI alternative could achieve the same purpose?
10. What promise or fear is accelerating adoption, and is it consistent with Christian hope?

References

- AI Singapore. 2026. "SEA-LION: Empowering Open Multilingual AI for Southeast Asia." <https://aisingapore.org/aiproducts/sea-lion/>.
- ASEAN Secretariat. 2024. *ASEAN Guide on AI Governance and Ethics*. https://asean.org/wp-content/uploads/2024/02/ASEAN-Guide-on-AI-Governance-and-Ethics_beautified_201223_v2.pdf.
- ASEAN Secretariat. 2025. *Expanded ASEAN Guide on AI Governance and Ethics: Generative AI*. <https://asean.org/book/expanded-asean-guide-on-ai-governance-and-ethics-generative-ai/>.
- Autio, Chloe, Reva Schwartz, Jesse Dunietz, Shomik Jain, Martin Stanley, Elham Tabassi, Patrick Hall, and Kamie Roberts. 2024. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>.
- Budd, Graham. 2023. "AI and Humanity: Navigating History's Next Great Wave." The Faraday Institute for Science and Religion, 30 November 2023.
- Cheng, Myra, et al. 2026. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." *Science*. <https://doi.org/10.1126/science.aec8352>.
- The Holy Bible, English Standard Version. 2016. Wheaton, IL: Crossway. Biblical references in this white paper are cited by book, chapter, and verse; wording is generally paraphrased rather than quoted.
- Department of Statistics Singapore. 2026. *General Household Survey 2025: Religion*. <https://www.singstat.gov.sg/publication-resources/general-household-survey-2025-religion>.
- Doshi, Anil R., and Oliver P. Hauser. 2024. "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content." *Science Advances* 10. <https://doi.org/10.1126/sciadv.adn5290>.
- Gmyrek, Pawel, et al. 2025. *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*. ILO Working Paper 140. <https://www.ilo.org/publications/generative-ai-and-jobs-refined-global-index-occupational-exposure>.
- Infocomm Media Development Authority. 2024. *Green Data Centre Roadmap*. <https://www.imda.gov.sg/how-we-can-help/green-dc-roadmap>.
- Infocomm Media Development Authority. 2026. *Model AI Governance Framework for Agentic AI*. <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>.
- International Energy Agency. 2025. *Energy and AI*. Paris: IEA. <https://www.iea.org/reports/energy-and-ai>.
- International Monetary Fund. 2024. "Impact of AI on Singapore's Labor Market." In *Singapore: Selected Issues*. IMF Country Report 24/256. <https://www.imf.org/-/media/files/publications/cr/2024/english/1sgpea2024002-print-pdf.pdf>.
- Langford, Michael D. 2022a. "A Theological Framework for Reflection on Artificial Intelligence." *SPU Works* 171. <https://digitalcommons.spu.edu/works/171>.
- Langford, Michael D. 2022b. "Artificial Intelligence and Theological Personhood." *SPU Works* 170. <https://digitalcommons.spu.edu/works/170>.
- Lennox, John C. 2024. *2084 and the AI Revolution: How Artificial Intelligence Informs Our Future*. Updated and expanded edition. Grand Rapids, MI: Zondervan Reflective. <https://zondervanacademic.com/products/2084-and-the-ai-revolution-updated-and-expanded-edition>.
- Leo XIV. 2026. *Magnifica Humanitas: On Safeguarding the Human Person in the Time of Artificial Intelligence*. Encyclical letter, 15 May 2026. <https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html>.

- Ministry of Manpower Singapore. 2026. "Written Answer to PQ on Strengthening Safeguards for Employment-Related AI." 5 May 2026. <https://www.mom.gov.sg/newsroom/parliament-questions-and-replies/2026/0505-written-answer-to-pq-on-strengthening-safeguards-for-employment-related-ai>.
- Lucky, Kate. 2023. "AI Will Shape Your Soul." *Christianity Today*, 11 September 2023. <https://www.christianitytoday.com/2023/09/artificial-intelligence-ai-chatgpt-spiritual-formation/>.
- Picard, Rosalind W. 2021. "The Future of Artificial Intelligence and Human Identity." *Faraday Paper* 24. The Faraday Institute for Science and Religion. <https://www.faraday.cam.ac.uk/wp-content/uploads/2021/11/Faraday-Paper-24-Picard.pdf>.
- Praxis. 2024. "A Redemptive Thesis for Artificial Intelligence." *The Praxis Journal*, 3 May 2024. <https://journal.praxislabs.org/a-redemptive-thesis-for-artificial-intelligence-ff7dafdd01b5>.
- Shoshani, Anat, Bar Gurfinkel, Ariel Kor, et al. 2026. "Attachment, Loneliness, and Social Support as Moderators of Conversational AI-Based Mental Health Outcomes." *npj Digital Medicine*. <https://doi.org/10.1038/s41746-026-02974-y>.
- Singapore Police Force. 2026. "Footage from Zoom Video Conference Involving Impersonation of Senior Government Officials." 16 May 2026. https://www.police.gov.sg/media-hub/news/2026/05/20260516_footage_from_zoom_video_conference_involving_impersonation_of_senior_government_officials.
- Skytland, Nicholas, Lauren Parsons, Alicia Llewellyn, Steele Billings, et al. 2026. "Evaluating Artificial Intelligence Through a Christian Understanding of Human Flourishing: The Flourishing AI Benchmark—Christian Single-Turn." arXiv:2604.03356. <https://arxiv.org/abs/2604.03356>.
- Smart Nation Singapore. 2026. "National AI Strategy." Last updated 21 May 2026. <https://www.smartnation.gov.sg/initiatives/national-ai-strategy/>.
- UNICEF Innocenti. 2025. *Guidance on AI and Children*, version 3.0. <https://www.unicef.org/innocenti/reports/policy-guidance-ai-children>.
- Williams-Ceci, Sterling, et al. 2026. "Biased AI Writing Assistants Shift Users' Attitudes on Societal Issues." *Science Advances*. <https://doi.org/10.1126/sciadv.adw5578>.

Note on interpretation

This white paper offers a broadly evangelical Christian framework. It draws appreciatively on Catholic social teaching and other Christian traditions without treating them as Protestant doctrinal authorities. Empirical claims remain subject to revision as evidence develops. The recommendations are principles for contextual judgment rather than a substitute for legal, clinical, cybersecurity, employment, or child-protection expertise.