
CHRISTIAN FAITH AND AI
Affirming Human Dignity amid Technological Advancement
Prof. Ir. Rudy C. Tarumingkeng, PhD
Professor of Management & AI Data Analyst
IBM-ASMI | RudyCT Academic Series
www.rudyct.com/cv.pdf
IMAGO DEI
Every person has inherent worth
NEIGHBOR-LOVE
Design for care and inclusion
JUSTICE & TRUTH
Challenge bias; verify AI outputs
STEWARDSHIP
Keep human responsibility and care for creation
Genesis 1:27 | Matthew 22:39 | Micah 6:8
AI should serve human flourishing, dignity, and the common good.

Christian Faith and AI: Affirming Human Dignity amid Technological Advancement

Prof. Ir. Rudy C. Tarumingkeng, PhD

www.rudyct.com/cv.pdf

IBM-ASMI | RudyCT Academic Series

Abstract

Artificial intelligence (AI) is changing how people learn, work, communicate, make decisions, and organize institutions. Its capacity to classify, predict, generate language, and automate tasks can serve human well-being, yet the same capacity can intensify surveillance, bias, exclusion, dependency, misinformation, and the reduction of persons to data profiles. This article develops a Christian ethical account of human dignity in the age of AI. It interprets the biblical affirmation that human beings are created in the image of God (imago Dei) in relation to embodiment, relationality, moral agency, vulnerability, work, justice, and care for creation. Scripture does not provide a technical manual for machine learning; it provides moral orientation for judging the purposes, means, and consequences of human action. Accordingly, AI is treated as a human-made instrument whose moral character depends on design, ownership, deployment, and governance. The article examines AI in education, employment, healthcare, public administration, media, church ministry, and environmental stewardship. It argues for a practical framework of dignity-centered discernment: define a legitimate human purpose; assess foreseeable harms and unequal effects; preserve meaningful human judgment and appeal; protect privacy and truth; distribute benefits fairly; and review outcomes with affected communities. A Christian response neither sanctifies technological progress nor rejects innovation by default. It asks whether technology serves neighbors, strengthens just institutions, and leaves persons more capable of truthful, responsible, and loving action. The article concludes with recommendations for churches, universities, businesses, governments, and Christian educators, with attention to Indonesia and the wider Global South.

Keywords

Artificial intelligence; Christian ethics; human dignity; *imago Dei*; human agency; justice; stewardship; responsible AI; church; Indonesia.

1. Introduction: A Human Question within a Technical Revolution

Artificial intelligence is often discussed through a language of speed, scale, accuracy, and productivity. These terms describe important properties of computational systems, but they do not answer the prior question: what ends should technology serve? An algorithm may produce a rapid answer, yet speed alone does not make the answer wise. A model may identify statistical patterns, yet a pattern is not automatically a morally legitimate basis for judging a person. A system may automate a task, yet automation does not determine who should be accountable when the task affects a life. The central Christian question is therefore not simply whether AI can perform a function. It is whether the function, its design, and its consequences honor the worth of persons and the common good.

The scale of current change makes this question urgent. Machine-learning systems now assist in image recognition, translation, medical triage, credit assessment, recruitment, customer service, content moderation, logistics, and public communication. Generative AI can produce text, images, code, audio, and synthetic video. These systems may help a small school prepare learning materials, help a hospital organize records, or help a local government translate information into regional languages. They may also fabricate evidence, reproduce discriminatory patterns, harvest personal information, displace workers without protection, and make consequential decisions difficult to challenge. The same technical capability can support care or facilitate harm, depending on the institutional setting and the power relations surrounding it.

Christian faith contributes a distinct but publicly discussable moral vocabulary. It speaks of creation, the image of God, neighbor-love, truth, justice, mercy, work, sin, responsibility, and hope. These concepts do not replace engineering, law, economics, or empirical evaluation. Rather, they help communities ask what a technical system is for, whom it benefits, whose labor and data it depends on, and who bears its errors. The biblical tradition also resists the temptation to confuse power with wisdom. In Genesis, the human vocation is not unlimited control but accountable care. In the prophets, prosperity is judged in relation to justice. In the Gospels, greatness is associated with service rather than domination.

This article argues that Christian faith can affirm technological creativity while insisting that the human person must never be reduced to a tool, score, prediction, or market segment. Human dignity is not earned through intelligence, productivity, independence, or digital competence. It belongs to persons because they are created and loved by God, embodied, relational, morally responsible, and open to communion with God and others. AI can assist human action, but it cannot take over the theological and moral responsibility of persons and institutions. The task is to build forms of innovation that enlarge human capacity for knowledge, care, participation, and justice.

2. Scope, Method, and Interpretive Approach

This is a narrative and interdisciplinary essay rather than an empirical study or a systematic review. It brings biblical interpretation into dialogue with Christian moral theology, contemporary AI ethics, and institutional practice. The biblical passages are read in their literary and theological contexts, not treated as isolated slogans or as direct predictions about present-day machines. The article draws especially on the creation narratives, the prophetic concern for justice, the wisdom tradition, the teaching of Jesus, and Pauline accounts of community and responsibility. It also engages current

discussions of responsible AI and recent ecclesial guidance, including the 2021 UNESCO Recommendation on the Ethics of Artificial Intelligence and the 2025 Vatican note *Antiqua et Nova*.

The term AI covers a diverse family of techniques, institutions, and products. In this essay it refers broadly to computational systems that perform tasks associated with learning, prediction, language, perception, planning, or decision support. Many contemporary systems are narrow and task-specific. Generative models produce plausible outputs from learned statistical patterns; this ability should not be confused with human understanding, conscience, or spiritual life. Some systems are deterministic, some probabilistic, and many combine software, data, human labor, and organizational procedures. Ethical analysis must therefore focus not only on a model in isolation but on the complete sociotechnical system: the data supply, design choices, procurement, deployment context, users, oversight, incentives, and routes for remedy.

The article uses a theological-ethical method of discernment. First, it identifies a biblical account of the human person and human vocation. Second, it examines where AI can support or threaten that account in ordinary institutional settings. Third, it translates moral commitments into governance questions and practices. The method is normative: it asks what ought to be done. It is also practical: it takes account of foreseeable effects, evidence, organizational constraints, and the voices of affected people. Biblical faith does not license technological claims unsupported by evidence. Nor should technical claims be treated as morally neutral simply because they are expressed in mathematical language.

The perspective is ecumenically Christian in its central commitments, while recognizing that churches differ in doctrine, ecclesial authority, and moral emphasis. The Vatican's *Antiqua et Nova* is cited as a significant contemporary Catholic contribution, not as a universal statement binding all Christians. The

article's arguments should be read as an invitation to dialogue among traditions and with people of other faiths and no religious affiliation. Dignity, justice, transparency, accountability, and protection from harm are also concepts in public ethics and human-rights frameworks. Christian reasoning can contribute to those common conversations without requiring public institutions to adopt specifically confessional premises.

3. What AI Can Do—and What It Does Not Establish

A system's apparent fluency can encourage people to attribute to it more than the available evidence warrants. A language model may generate a coherent answer, imitate a style, summarize a document, or respond with conversational warmth. These achievements arise from computational processes trained on data and shaped by design and deployment. The system can be useful without possessing the human experiences that words often express. Producing sentences about grief is not the same as grieving; generating a prayer is not the same as praying; identifying a moral argument is not the same as having a conscience that is answerable to God and neighbor.

This distinction is not a claim that machines are insignificant. AI can have substantial effects because people rely on it, organizations embed it in decisions, and institutions grant it authority. A system need not be conscious to shape access to employment, credit, healthcare, education, or public services. Indeed, the danger may be greatest when powerful systems are treated as neutral or inevitable. Their outputs can acquire institutional force even when the underlying model is uncertain, biased, or poorly suited to the context. Moral significance therefore concerns both the system's actual properties and the human practices through which it is used.

Christian theology has historically treated intelligence as more than calculation. Human intelligence includes bodily perception, memory, emotion,

practical judgment, moral deliberation, imagination, social learning, and the ability to seek meaning. It develops through relationships, language, culture, education, work, and spiritual formation. Human beings can revise their purposes, confess wrongdoing, forgive, and respond to another person as a neighbor rather than as a data point. These features do not imply that every human decision is wise or just. They do remind us that intelligence is embedded in a life and a community, not merely in the production of correct-looking outputs.

The 2025 note *Antiqua et Nova* distinguishes the instrumental capacities of AI from human intelligence understood in its embodied and relational dimensions (Dicastery for the Doctrine of the Faith & Dicastery for Culture and Education, 2025). This distinction is especially relevant when an organization treats a model's score as an authoritative judgment. A risk estimate may inform a clinician, but it does not itself care for a patient. A hiring model may rank applications, but it cannot assume the employer's duty to provide fair consideration. A chatbot may help someone find a Bible passage, but it cannot become a pastor or a community of faith. The proper response is neither to deny useful machine capabilities nor to project human personhood onto them. It is to assign authority and responsibility carefully.

4. The Biblical Foundation of Human Dignity

Genesis 1:26–27 declares that humankind is created in the image and likeness of God. The passage gives no technical definition of the image, and Christian traditions have interpreted it in several ways: as a capacity for reason, a vocation of stewardship, a relational reality, a status bestowed by God, or a combination of these. For the purposes of AI ethics, a crucial implication is that the worth of a person is not contingent on performance. If dignity depended on superior calculation, memory, speed, or productivity, machines could appear to threaten human worth whenever they surpassed particular

human capacities. The biblical affirmation points elsewhere: the human person has value before any system measures or ranks that person.

The image of God belongs to humanity as a shared and inclusive status. Genesis names male and female together (Genesis 1:27); the text does not reserve dignity for rulers, experts, able-bodied people, or economically productive citizens. The Bible's broader witness extends moral concern to strangers, widows, orphans, the poor, the sick, and those without power. Deuteronomy 10:17–19 joins God's impartiality with a command to love the stranger. James 2:1–9 condemns partiality toward the wealthy. These passages challenge AI systems that reproduce unequal treatment or make exclusion harder to see. If an automated system disproportionately denies services to a marginalized group, its technical sophistication cannot excuse the institution from examining the harm.

Dignity also has a relational dimension. Human beings are not isolated units but persons whose lives are formed and sustained in relationships—with God, family, community, institutions, and the created world. The biblical story describes human flourishing through covenant, mutual responsibility, and belonging. The command to love one's neighbor (Leviticus 19:18; Mark 12:29–31) places the other person before us as someone to be encountered, not merely predicted. This does not prohibit all classification; institutions need categories to organize services. It does require that categories remain subordinate to attentive human care and that a person be able to explain, contest, and sometimes transcend a profile assigned by a system.

Human dignity is embodied. The biblical account of creation describes human life as formed from the earth and animated by God's breath (Genesis 2:7). The incarnation, in Christian confession, further affirms the significance of embodied human life (John 1:14). A digital representation can capture information about a person, but it cannot contain the whole person. Health records, click histories, facial images, voice samples, and behavioral

predictions are partial traces. When institutions treat such traces as a complete identity, they risk turning people into objects of surveillance and control. Respect for dignity requires data minimization, contextual integrity, meaningful consent where possible, and limits on secondary use.

The Bible also recognizes vulnerability and dependence. Human worth does not disappear through illness, age, disability, poverty, error, or the need for assistance. The parable of the Good Samaritan (Luke 10:25–37) makes neighbor-love concrete through costly attention to a wounded stranger. A dignity-centered system should therefore be judged by how it treats people who have the least power to refuse, understand, or correct it. A tool that improves average efficiency while abandoning those who cannot navigate a digital interface may be administratively efficient and morally deficient. The biblical measure is not only the success of the capable but the treatment of the vulnerable.

5. Human Vocation, Stewardship, and the Limits of Dominion

Genesis 1:28 and 2:15 portray human beings as entrusted with the earth and placed in a garden "to till it and keep it" (NRSVUE). Christian traditions have sometimes interpreted dominion as permission for unrestricted exploitation. A more careful reading holds authority together with responsibility: the created world belongs to God, and human beings are accountable stewards. This matters for AI because technological capacity can encourage a culture of mastery. Data can be collected simply because it can be collected; behavior can be manipulated because it can be predicted; labor can be automated because it can be made cheaper. Stewardship asks a different question: what forms of use protect life, sustain communities, and preserve the conditions of flourishing?

Stewardship is not passive conservation. It includes creative work, scientific inquiry, and the use of tools to address human needs. AI may help identify

crop disease, forecast flood risks, improve energy efficiency, translate health guidance, or make public information more accessible. Such applications can be expressions of responsible care when designed with local knowledge and evaluated against real outcomes. In a country exposed to earthquakes, floods, volcanic hazards, and unequal access to services, predictive systems may help authorities prepare and allocate resources. Yet the same systems can fail when data omit remote communities, when warnings are inaccessible to people with disabilities, or when officials mistake a forecast for certainty.

The biblical account also limits human claims to ultimate authority. The prohibition against idolatry (Exodus 20:3–6) is not only about statues; it warns against giving ultimate trust to finite things. AI becomes an idol when people treat it as an all-knowing oracle, a substitute for moral judgment, or a source of salvation from uncertainty and conflict. This warning is not a reason to avoid computational tools. It is a reason to keep them in their proper place. A model can illuminate a decision, but it cannot decide what is worthy of love, what justice requires, or what sacrifices a community should accept.

The story of Babel (Genesis 11:1–9) offers another caution. Its concern is not that human beings cooperate or build, but that collective power can become self-aggrandizing and resistant to limitation. Contemporary AI development can similarly concentrate technical and economic power in a small number of firms, states, or institutions. When decisions about data, infrastructure, language, and platform rules are made without public participation, communities become dependent on systems they did not shape. A Christian account of stewardship therefore includes institutional restraint, shared governance, and accountability for those who control powerful technologies.

6. Technology, Sin, and Moral Responsibility

Christian ethics does not assume that good intentions automatically produce good outcomes. The doctrine of sin describes the ways human beings and

institutions distort their capacities, relationships, and purposes. In technological settings, sin may appear as greed that prioritizes extraction over care; pride that treats prediction as omniscience; fear that encourages surveillance; indifference that externalizes harm; or self-deception that hides value choices behind technical vocabulary. These patterns do not make technology inherently evil. They explain why moral reflection must examine the incentives and power structures surrounding its design and use.

AI systems are made and maintained through human decisions. People choose what problem to address, what data to collect, which outcomes to optimize, what errors to tolerate, which groups to test, and what level of human review to provide. The resulting model may behave in ways its creators did not fully anticipate, but unpredictability does not erase responsibility. It changes what responsible conduct requires: testing, monitoring, documentation, incident response, and willingness to suspend a system when evidence indicates unacceptable harm. Responsibility should be distributed according to actual roles and power, while remaining traceable enough that no organization can claim that "the algorithm" alone made the decision.

The New Testament's moral vision places freedom within accountability. Galatians 5:13 says that freedom is expressed through service in love. Romans 14:12 reminds believers that each person will give an account to God. These passages resist both technological fatalism ("innovation cannot be stopped") and simplistic voluntarism ("each user can solve the problem alone"). Developers, executives, procurement officers, regulators, professionals, and users have different responsibilities. The organization that deploys an AI hiring system has greater control over its consequences than an applicant who must submit to it. Ethical governance should assign duties in proportion to power, knowledge, and capacity to prevent harm.

Confession and repentance also have practical analogues. An institution should be able to acknowledge when an AI system caused harm, tell affected people what happened, provide remedy, and change the system or its use. A public statement of principles is inadequate if the organization refuses to correct a harmful practice. Christian repentance is more than regret; it involves turning from wrongdoing and repairing relationships where possible. In organizational life this can mean withdrawing an unfit model, reopening decisions, compensating losses, revising procurement standards, and listening to people whose experiences were previously ignored.

7. Neighbor-Love as a Test of Design and Use

Jesus summarizes the law through love of God and love of neighbor (Matthew 22:37–40). Neighbor-love is not a sentimental feeling added after technical decisions are complete. It provides a test for the whole life cycle of AI. Who is treated as the neighbor in the design brief? Whose needs are made visible in the data? Who has the opportunity to explain a decision? Who receives the benefits and who absorbs the cost? Does the system make it easier to respond to people with patience, or does it encourage institutions to avoid them? These questions convert a theological virtue into a practical discipline of attention.

The Good Samaritan narrative is especially instructive for automated service systems. The Samaritan sees a person in need, interrupts his journey, approaches the injured man, and assumes material responsibility for care (Luke 10:30–37). The story does not imply that every service must be delivered face-to-face or that automation is incompatible with compassion. It does show that moral responsibility requires more than categorizing a case. A hospital chatbot may help patients find information, but there must be a human pathway when a person is frightened, confused, or at risk. A government portal may speed up applications, but a citizen who cannot use it must not be abandoned to a digital dead end.

Neighbor-love also includes truthful communication. Ephesians 4:25 calls the community to speak truthfully, and Proverbs repeatedly warns against false witness. Generative AI makes it easier to fabricate images, voices, documents, and persuasive text at scale. Such material can damage reputations, manipulate voters, defraud older adults, or inflame religious and ethnic tensions. Christian communities should develop habits of verification before sharing content, especially when a message confirms existing fears or attacks a vulnerable group. Platforms and publishers also bear responsibility to label synthetic media where appropriate, reduce deceptive amplification, and offer accessible routes for correction.

Love does not mean avoiding disagreement or withholding legitimate criticism. It means that a person should not be exposed to deception, harassment, or unjust exclusion for the convenience of an institution. A Christian framework therefore joins care with truth, and compassion with accountability. When a model is used in a consequential context, affected individuals should receive understandable information, a meaningful opportunity to contest an outcome, and access to a human decision-maker able to change the result. A nominal "human in the loop" is insufficient if the human reviewer lacks time, authority, or relevant knowledge. The person who seeks remedy must encounter an accountable human process, not a symbolic signature.

8. Justice, Bias, and Structural Inequality

The biblical prophets repeatedly judge societies by their treatment of the poor and their use of power. Amos condemns those who "trample on the poor" and manipulate economic institutions (Amos 5:11–12); Micah calls for justice, kindness, and humility (Micah 6:8). These texts invite analysis not only of individual prejudice but of structures that distribute opportunity and risk. AI systems trained on historical records may reproduce patterns shaped by past discrimination. They may also create new forms of exclusion when data

are incomplete, proxies are misleading, or design teams assume that the experiences of a dominant group are universal.

Bias is not one technical defect with one technical fix. It can enter through the definition of a problem, sampling choices, labels, measurement, model objectives, deployment thresholds, and institutional responses to errors. A recruitment model trained on past hiring data may learn to favor profiles associated with previous employees. A credit system may penalize neighborhoods because historical disadvantage is encoded in financial records. A facial recognition system may have different error rates across demographic groups. Fairness testing can reveal some disparities, but it cannot by itself determine whether the original decision rule is ethically justified. Sometimes the just response is not to improve the prediction but to stop using the system for that decision.

Christian justice asks whether social arrangements enable people to participate in life and to receive what is due to them. It therefore includes distributive questions: who owns the infrastructure, who profits from data, which communities receive investment, and whether productivity gains are shared. It also includes procedural justice: can a person understand and challenge a decision? Recognition justice asks whether the system represents groups respectfully and accommodates language, disability, and cultural context. Reparative justice asks how to address injuries already caused. These dimensions overlap with, but are not reducible to, statistical parity metrics.

Indonesia's social and linguistic diversity makes contextual fairness especially important. A system trained largely on high-resource languages and urban data may misunderstand local expressions, accents, informal work, or community-specific needs. Digital services can increase access in some regions while excluding people with limited connectivity, older devices, low digital literacy, or disabilities. Christian institutions should not assume that a product marketed as “inclusive” has been tested with Indonesian

communities. They can ask for evidence about training data, subgroup performance, accessibility, data protection, local language quality, and mechanisms for appeal. They can also involve local pastors, teachers, health workers, community organizations, and users in evaluating whether an application is helpful and safe.

Justice further requires attention to the global distribution of AI's material costs. Data centers require electricity, water, minerals, and land. Content moderation and data labeling may rely on low-paid labor, including workers exposed to disturbing material. Communities that host infrastructure may bear environmental burdens without receiving commensurate benefits. A Christian ethic of solidarity asks technology companies and institutions to disclose these dependencies and to improve labor and environmental conditions throughout the supply chain. Human dignity includes the people whose invisible work makes apparently effortless automation possible.

9. Truth, Knowledge, and the Risk of False Authority

Christian faith has a deep concern for truth. In John 8:32, truth is associated with freedom; in John 14:6, truth is central to Jesus' self-description. These passages do not equate truth with the output of any institution or person. They do insist that deception is spiritually and socially serious. AI-generated answers can be fluent and wrong, citing nonexistent sources or combining accurate and inaccurate claims. Users may over-trust the output because it sounds confident. In academic, legal, medical, and pastoral settings, a plausible error can cause lasting harm.

The ethical issue is not merely “hallucination” as a technical term. It is the human practice of presenting unverified output as knowledge. A researcher remains responsible for checking sources and interpreting evidence. A student remains responsible for the claims submitted under the student's name, within the rules of the institution. A pastor remains responsible for

doctrinal and pastoral counsel. A company remains responsible for customer-facing claims. AI can support inquiry by suggesting search terms, organizing notes, or offering a draft, but the user must verify facts, trace references, and disclose use when disclosure is required.

The difference between information and wisdom is especially important. Information can be retrieved or generated; wisdom involves judging what is good, fitting, and responsible in a particular situation. Proverbs 4:7 commends the pursuit of wisdom, while James 1:5 invites believers to ask God for wisdom. An AI assistant may present options, but it cannot bear the spiritual discipline of discernment. It cannot know the full significance of a family conflict, a congregational wound, or a patient's fears merely from the data entered. Wisdom requires listening, humility, prayer, experience, and accountability within community.

Misinformation also harms public life. Synthetic political content, fake endorsements, fabricated sermons, and manipulated audio can undermine trust. Christian leaders have a responsibility to model careful speech and to correct falsehoods without amplifying them unnecessarily. Congregations can teach members to examine source, date, context, and evidence; to distinguish fact from interpretation; and to avoid forwarding emotionally charged content before verification. At institutional level, media literacy and provenance practices should accompany technical tools. A church or university can provide a designated channel for reporting synthetic impersonation and clarify how official communications are authenticated.

10. Privacy, Surveillance, and the Integrity of the Person

Privacy is more than secrecy. It protects the space in which people form relationships, seek counsel, explore belief, and make choices without unjust observation. The biblical tradition includes the importance of trust, confession, and discretion. Proverbs 11:13 praises the trustworthy person who

keeps confidence; Matthew 6:6 refers to prayer in a private place. These texts do not prescribe a modern data policy, but they remind communities that not every personal fact should be exposed or converted into an institutional asset.

AI increases the capacity to infer sensitive information from ordinary data. Location patterns, browsing behavior, voice characteristics, facial images, and social networks may reveal health conditions, political affiliation, religious practice, or emotional vulnerability. A person may not have knowingly consented to such inferences. In churches, schools, and workplaces, data collection can be particularly coercive because people may fear losing access or approval if they refuse. Institutions should collect only what is needed for a legitimate purpose, define retention periods, restrict access, secure data, and prohibit incompatible secondary use. Consent should be specific and meaningful rather than hidden in a long document or bundled with essential services.

Surveillance can also alter behavior. When people know that every interaction may be monitored, they may avoid asking questions, seeking help, or expressing dissent. The result is not only loss of privacy but a weakening of trust and agency. A Christian community should be cautious about using AI to infer spiritual commitment, rank members, monitor attendance patterns for punitive purposes, or predict who may leave a congregation. Such practices may treat pastoral care as population control. Data can help identify unmet needs, but care must remain voluntary, relational, and proportionate.

Privacy protections should include special care for children, people in crisis, migrants, and others with limited power. Institutions must also consider cybersecurity and the risks of data breach. A data set intended for ministry or education can be repurposed or exposed. Leaders should ask: What would happen if this information became public? Would the person reasonably expect it to be used this way? Is there a less intrusive means to achieve the same goal? Who can access the data, and who is accountable if it is misused?

These questions translate respect for the person into concrete stewardship of information.

11. Work, Automation, and Economic Dignity

Work is a significant part of human life, though it does not define a person's worth. Genesis portrays work as part of human vocation, and the Bible affirms the value of labor while criticizing exploitation. The command to rest in Exodus 20:8–11 places limits on production and recognizes that human life is more than work. AI may relieve people of dangerous, repetitive, or burdensome tasks. It may also threaten livelihoods, intensify work through monitoring, devalue professional expertise, or shift economic gains toward owners of capital and data.

A Christian assessment must avoid two errors. The first is to romanticize every existing job, including work that is degrading or unsafe. The second is to celebrate productivity while ignoring workers who lose income, status, bargaining power, and community. The proper question is how technological change can be organized so that it supports dignified work and a fair transition. Employers should assess likely effects before deployment, consult workers and unions, provide retraining tied to actual opportunities, and share productivity gains where possible. Governments can support income security, lifelong learning, and labor protections that keep pace with new forms of work.

The parable of the laborers in the vineyard (Matthew 20:1–16) does not provide a modern wage formula, but it challenges assumptions that market productivity is the sole measure of deservingness. Christian communities can therefore question economic models that treat workers as disposable inputs. If an AI system removes tasks from a role, the organization should ask what new responsibilities emerge, whether the remaining work becomes more intense, and how workers can participate in redesign. Workplaces should not

use algorithmic management to set impossible pace, penalize legitimate breaks, or make disciplinary decisions opaque.

In Indonesia, the consequences of automation will vary across formal employment, informal work, small enterprises, public services, and regional economies. A large business may absorb transition costs that a small enterprise cannot. A digital platform may provide income access but also shift risk to contractors. Universities and churches can help workers understand AI tools without implying that every worker must become a programmer. Digital literacy, domain expertise, communication, ethical judgment, and the capacity to learn remain valuable. A just transition requires specific pathways, not slogans about reskilling.

12. AI in Education and the Formation of Character

Education is not merely the transfer of information. It forms judgment, habits, relationships, imagination, and responsibility. Deuteronomy 6:6–7 presents teaching as embedded in daily life, while the Gospels portray learning through dialogue, example, questioning, and community. AI tutors, translation tools, accessibility systems, and writing assistants can support learners. They may help a student receive explanations at different levels, practice a language, or obtain feedback outside classroom hours. They can also encourage dependence, facilitate plagiarism, expose sensitive student data, or narrow learning to what can be easily measured.

A Christian educational approach asks whether AI strengthens the student's capacity to think, create, and act responsibly. Teachers can use generative systems to develop examples, generate practice questions, or adapt materials, while making their own pedagogical judgment visible. Students can use AI to brainstorm or receive feedback when course policies permit, but should learn to verify claims, cite sources, disclose assistance, and distinguish their own reasoning from generated text. Assessment should value process, oral

explanation, reflection, collaboration, and application, not only polished submissions that a model can produce.

The risk of unequal access matters. Students with paid tools, strong connectivity, and confident digital skills may gain advantages over those without them. Institutions should not require students to submit personal data to commercial platforms without careful review. They should ensure that AI-supported instruction remains accessible to students with disabilities and responsive to local languages and contexts. Human teachers remain essential because education involves recognition, encouragement, ethical example, and responsiveness to a student's circumstances. A model can offer feedback, but a teacher can notice that a learner is discouraged, confused, or carrying burdens not visible in a data record.

For Christian schools and universities, AI literacy should include theological and ethical reflection without turning technology lessons into fear-based instruction. Students can study how training data shape outputs, how recommendations influence attention, and how automated decisions distribute benefits. They can compare a system's claims with biblical principles of truth, justice, dignity, and service. Such learning prepares graduates to participate in technological change rather than merely consume it. The educational aim is not to compete with machines at speed but to become more capable of judgment, collaboration, and responsible creativity.

13. Healthcare, Care, and the Limits of Automation

AI can support medical imaging, triage, documentation, drug discovery, and public-health surveillance. In settings where specialists are scarce, carefully validated systems may help clinicians identify risks or prioritize attention. Yet healthcare is not only a classification problem. It includes bodily suffering, uncertainty, consent, trust, and decisions about what burdens a patient is willing to bear. Jesus' healing ministry repeatedly attends to people as

persons, restoring not only bodily function but also social inclusion (for example, Mark 1:40–45; Luke 8:43–48).

An AI system should support clinicians rather than obscure clinical responsibility. A recommendation is useful only if the clinician understands its limitations and can override it. Patients should know when AI materially contributes to care, what data are used, and how to raise concerns. Systems require validation in populations and clinical environments similar to those in which they will be used. A model trained on data from one country or demographic group may not perform equally elsewhere. When errors occur, responsibility should not be deflected to the patient or to an abstract algorithm; hospitals, vendors, and professionals must have clear duties.

Healthcare data are especially sensitive. Health systems should protect confidentiality and avoid using patient data for unrelated commercial purposes without a lawful and ethically justified basis. Predictive analytics can be used to allocate resources, but must not become a mechanism for denying care to patients deemed costly or unlikely to benefit. A Christian ethic of dignity requires special attention to the sick, disabled, elderly, and poor, whose interests may be weakly represented in institutional data. A technically accurate forecast can still be used unjustly if the institution's objective is morally wrong.

AI can also make care more humane when it reduces administrative burden and gives professionals more time for patients. This outcome is not automatic. If efficiency savings lead to higher caseloads and fewer conversations, patients may experience less care despite more sophisticated tools. Leaders should evaluate not only throughput and cost but also patient understanding, continuity, dignity, access, and trust. In healthcare, the measure of a tool includes whether it enables clinicians to be more present to those who suffer.

14. Public Administration, Power, and the Common Good

Public institutions use data systems to deliver benefits, detect fraud, plan transport, manage emergencies, and allocate services. These uses can improve coordination and reduce delays. They can also create opaque eligibility decisions, expand surveillance, and make citizens dependent on systems that are difficult to challenge. Romans 13:1–7 is sometimes read as an unconditional endorsement of authority; the broader biblical tradition also holds rulers accountable to justice. The prophets confront unjust power, and Acts 5:29 recognizes limits to obedience when human authority conflicts with higher moral duty.

A democratic and dignity-centered administration should identify high-impact uses, publish understandable information about them, conduct independent assessments, and provide appeal routes. People affected by a decision should be able to know the relevant reasons and correct inaccurate data. Public agencies should not outsource accountability to vendors or treat proprietary secrecy as a blanket excuse. Procurement contracts should specify audit access, data governance, security, performance expectations, and remedies. A system that cannot be examined sufficiently should not be granted authority over consequential public decisions.

AI in public services should be evaluated against the common good, a concept shared by Christian social thought and many civic traditions. The common good is not merely the sum of individual preferences; it concerns the social conditions that allow people and communities to flourish. This includes safety, education, healthcare, participation, accessible information, and environmental sustainability. AI may help public agencies coordinate resources during floods or epidemics, but the benefits depend on reliable infrastructure, clear authority, local capacity, and trust. A model cannot compensate for underfunded services, poor data stewardship, or political unwillingness to serve marginalized communities.

The Indonesian context calls for careful attention to national and local capacity. Digital transformation should not mean importing systems that cannot be maintained, audited, or adapted by Indonesian institutions. Public agencies should build technical capacity, include local universities and civil society, and strengthen privacy and accountability mechanisms. They should also ensure that citizens who prefer offline or assisted channels are not excluded. A state that uses technology in the name of efficiency must still remain reachable by the people it serves.

15. AI, Media, and the Integrity of Christian Witness

Christian witness is communicated through speech, conduct, worship, teaching, and service. Digital media can extend access to sermons, Bible study, and pastoral resources. Translation and transcription tools may help people who speak different languages or live with disabilities. AI can assist a small congregation with routine communications, but it can also produce inaccurate theological explanations, imitate a pastor's voice, or create false announcements. If churches use automated content, they should identify who reviews it and who is responsible for its pastoral and doctrinal claims.

The temptation to use AI for scale can neglect the relational nature of ministry. A congregation is not simply an audience to be optimized for engagement. Pastoral care involves listening, confidentiality, discernment, and presence. A chatbot can direct a person to a resource or provide basic information, but it should not be presented as a substitute for a pastor, counselor, or emergency service. Clear boundaries are essential when someone discloses abuse, suicidal thoughts, domestic violence, or a serious medical crisis. The system should direct the person to appropriate human support and avoid creating a false impression of confidentiality or competence.

The command to love one another (John 13:34–35) calls churches to treat digital communication as part of moral life. Congregations can establish rules for the use of synthetic images, voice, and text; verify public statements; obtain permission before processing personal stories; and protect children's information. Church leaders should not use AI to simulate testimonials, fabricate attendance, or manipulate giving. The means of evangelization matter because they shape the credibility of the message. Truthful and transparent communication is part of Christian witness.

AI may also support theological reflection by helping readers compare translations, organize historical sources, or locate themes. These uses should be accompanied by scholarship, context, and humility. A model may summarize a passage in a way that misses genre, historical setting, or doctrinal complexity. Pastors and teachers should show how interpretation is tested against Scripture, tradition, reason, and the life of the community. The tool can assist the study of theology; it cannot replace the community's responsibility to interpret and live the faith.

16. Human Agency, Accountability, and Meaningful Oversight

Human agency includes the ability to understand reasons, form intentions, deliberate, act, and answer for one's conduct. AI can extend agency by helping people overcome barriers, explore options, and carry out tasks. It can diminish agency when it narrows choices, manipulates attention, conceals decision rules, or makes refusal practically impossible. The ethical goal is not to keep humans involved in every computation. It is to ensure that people retain meaningful authority over decisions that affect rights, safety, livelihood, and relationships.

“Human oversight” is often invoked as a safeguard, but its quality varies. A reviewer may be expected to approve hundreds of automated recommendations quickly, lack access to the data or model limitations, or be

penalized for disagreeing with the system. Such oversight is ceremonial. Meaningful oversight requires sufficient time, competence, authority, access to relevant information, and protection for dissent. It also requires clarity about when the human must intervene and how the institution learns from errors. If a decision is too complex or harmful to explain and contest, the organization should reconsider whether it should be automated at all.

Accountability should follow the chain of decisions. Developers have duties concerning design, testing, documentation, and known limitations. Vendors have duties concerning claims, updates, security, and support. Deploying institutions have duties to assess context, train staff, inform affected people, monitor outcomes, and provide remedy. Professionals have duties to exercise judgment within their expertise. Regulators have duties to set enforceable standards and investigate serious failures. Users have responsibilities too, but those responsibilities must not be used to shift institutional obligations onto people with little power.

Practical accountability includes an inventory of AI systems, risk classification, impact assessments, documentation, incident reporting, redress mechanisms, and periodic review. It also requires a culture in which employees can raise concerns without retaliation. The biblical tradition values truth-telling, correction, and restoration. A Christian organization should therefore make it possible to question a system before harm becomes normalized. Accountability is not hostility toward innovation; it is a condition for trustworthy innovation.

17. Creation Care and the Material Footprint of AI

Digital technology can appear immaterial because users encounter it through screens and “cloud” services. In reality, AI depends on data centers, electricity, cooling water, semiconductors, minerals, transport, and human labor. Training and running large models can have substantial environmental costs, though

those costs vary by system, energy source, utilization, and measurement method. Christian stewardship therefore asks organizations to consider the full life cycle of AI, including energy efficiency, hardware use, water consumption, e-waste, and supply-chain conditions.

Creation care does not require rejecting every energy-intensive system. It requires comparing benefits with burdens and considering less resource-intensive alternatives. A smaller model may be sufficient for a local task. A system may be deployed only where it offers meaningful benefit. Organizations can request environmental reporting from providers, improve data-center efficiency, and avoid duplicative computations. Public institutions should consider whether AI investment competes with essential infrastructure or whether it enables substantial environmental gains, such as improved grid management or disaster forecasting.

Romans 8:19–23 portrays creation as groaning and awaiting liberation. The passage is not a policy document, but it resonates with the responsibility to address ecological harm. AI can assist conservation and climate adaptation, yet it can also accelerate extraction, consumption, and surveillance. An ethical assessment must include both sides. Communities affected by mining or infrastructure development should have a voice in decisions. Environmental benefits should not be claimed without evidence, and claims of “green AI” should be open to independent scrutiny.

For Indonesia, environmental evaluation should account for the energy mix, water stress, geographic distribution of infrastructure, disaster exposure, and local communities. Universities and churches can help connect digital ethics with ecological literacy. The aim is a form of innovation that respects both human dignity and the integrity of the natural world, recognizing that human well-being depends on ecological systems that technology cannot replace.

18. A Dignity-Centered Framework for Christian Discernment

A practical framework can help Christian institutions move from broad principles to decisions. The first question is purpose: what human need does the system address, and is AI an appropriate means? The organization should articulate a specific public or pastoral benefit rather than adopting a tool because it is fashionable. It should consider alternatives, including improving a process, hiring staff, or using a simpler technology. The second question concerns affected persons: who will be represented in the data, who may be excluded, and who will experience the consequences? The people most affected should be consulted early enough to influence design and procurement.

The third question is harm and justice. Institutions should identify foreseeable risks to privacy, safety, access, employment, reputation, and equality. They should examine performance across relevant groups, evaluate environmental and labor effects, and determine whether the system may intensify existing disadvantage. The fourth question is agency: can people understand the decision, refuse nonessential data collection, obtain human help, and appeal an adverse outcome? The fifth is accountability: which named role owns the system, monitors it, records incidents, and provides remedy? Responsibility must be assigned before deployment, not improvised after a failure.

The sixth question is truth and competence. Are the system's limits known? Are outputs verified before they are used in high-stakes settings? Are staff trained to detect error and automation bias? The seventh question is proportionality: does the expected benefit justify the data collected, degree of automation, environmental cost, and possible harm? A system with low public benefit and high intrusion should be rejected or redesigned. The eighth question concerns review: when will the organization reassess the system, and what evidence would lead it to modify, pause, or retire it?

These questions can be summarized in an operational cycle: purpose, participation, assessment, safeguards, accountable use, and review. The cycle aligns with responsible-AI frameworks emphasizing human rights, transparency, fairness, privacy, and accountability (Floridi et al., 2018; Jobin et al., 2019; UNESCO, 2021). Its Christian contribution is to root these practices in the conviction that every person is a neighbor and that power is entrusted for service. Such a framework should not become a checklist that substitutes for judgment. It is a discipline for asking better questions and documenting responsible choices.

19. Applications in Churches, Universities, Business, and Government

Churches can begin with an inventory of AI uses: sermon preparation, translation, member databases, livestream moderation, donor systems, counseling resources, and communications. They can adopt a short policy that identifies acceptable and unacceptable uses, requires human review for public theological content, protects confidential information, and prohibits synthetic impersonation without clear consent. Leaders can provide training for staff and congregants on misinformation, privacy, and responsible use. Consultation with people who receive pastoral care can reveal risks that leaders may not see from an administrative perspective.

Universities should combine AI literacy with academic integrity and student protection. They should define transparent course-level expectations, provide secure and equitable access where tools are required, and avoid making detection software the sole basis for misconduct judgments. AI-detection systems can produce false positives and should not replace a fair process. Faculty can redesign assessments to make reasoning visible and to teach citation and verification. Research ethics committees should examine AI-related data use, model validation, authorship, and participant consent. Universities can also establish interdisciplinary forums where computer

scientists, theologians, social scientists, and affected communities deliberate together.

Businesses can conduct impact assessments before procurement, include worker and customer perspectives, and require vendors to document data provenance, model limitations, security, and environmental impacts. Human resource departments should be especially careful with recruitment, evaluation, scheduling, and dismissal systems. Employees should receive notice and have a channel to contest automated outcomes. Companies can measure the distribution of productivity gains and create transition plans where roles change. Ethical commitments should be tied to executive responsibility, budgets, and audits rather than public-relations language alone.

Governments should set clear standards for high-risk uses, preserve due process, enable independent oversight, and invest in public capacity. They should avoid overreliance on vendors for essential services and ensure that citizens can access services through human and offline channels. Policy should address data protection, cybersecurity, discrimination, labor transition, and environmental impact together. International cooperation is needed because AI systems, supply chains, and data flows cross borders. Global standards can provide a baseline, while local implementation must account for social context, language, institutional capacity, and public participation.

20. Indonesia and the Global South: Context, Capacity, and Voice

The Global South is not a single context, and Indonesia itself contains diverse regions, languages, religions, economies, and levels of infrastructure. The ethical challenge is not merely to gain access to AI products developed elsewhere. It is to shape the terms on which technology is adopted, governed, and adapted. This includes the ability to evaluate vendors, secure data, audit outcomes, develop local expertise, and withdraw from systems that do not

serve the public interest. Digital sovereignty should not become isolationism; it should mean meaningful capacity and accountable choice.

AI may help address real needs: agricultural extension, disaster preparedness, translation, accessible education, public-health information, and support for small businesses. These applications are most promising when co-designed with communities and tested in real conditions. Local knowledge should not be treated as noise to be cleaned from a data set. A flood-risk tool, for example, may benefit from combining remote sensing with community observations about drainage, informal settlements, and evacuation routes. A language model that handles Bahasa Indonesia but not local languages may have limited value for people who need it most.

Risks include dependence on foreign infrastructure, extractive data practices, unequal access, low-quality automated content, and decisions made from incomplete records. The power to define categories can be as consequential as the power to own hardware. If a system treats informal work as evidence of unreliability, or interprets local linguistic patterns as error, it can reproduce injustice under a veneer of objectivity. Indonesian universities, public agencies, civil society, religious organizations, and private firms can jointly develop standards for local evaluation, open educational resources, and public-interest use cases.

Christian institutions can contribute through education, social service, health care, and community networks. They can convene dialogue across faiths and sectors, identify harms experienced by vulnerable groups, and support digital literacy. Their contribution will be credible when they practice what they teach: protect member data, correct misinformation, disclose AI assistance where relevant, and prioritize human care over efficiency metrics. In this way, Christian ethics can participate in a plural public conversation about technology and the common good.

21. Objections and Replies

One objection is that theology should remain outside AI governance because public policy must be secular and pluralistic. The concern is understandable: no religious institution should impose doctrine on citizens who do not share it. However, public reasoning can receive contributions from religious traditions when arguments are expressed in accessible terms and remain open to scrutiny. Christian claims about human dignity, care for the vulnerable, and limits on power can be discussed alongside human-rights, democratic, and professional frameworks. The article does not propose that public institutions enforce a Christian theology; it proposes that Christian communities act responsibly and contribute to shared ethical deliberation.

A second objection is that AI is simply a tool, and moral responsibility belongs only to users. This is partly right: machines do not bear responsibility in the same way that persons and institutions do. But the phrase “just a tool” can conceal design choices and power relations. A tool's affordances shape action, and an organization that selects and deploys it creates conditions for predictable use. A hiring system's impact cannot be assigned solely to an applicant who never chose it. Responsibility belongs to the people and institutions that design, procure, configure, and govern the tool, in proportion to their knowledge and control.

A third objection is that human judgment is itself biased, inconsistent, and sometimes cruel. This is true. Human involvement is not a magical guarantee of justice, and automation may improve consistency in some tasks. The relevant comparison is not between imperfect machines and ideal humans; it is between actual sociotechnical arrangements. Human decisions can be audited and appealed; automated decisions can too, if institutions build those mechanisms. The goal is to combine computational strengths with accountable human judgment, while preventing either from escaping scrutiny.

A fourth objection is that strong safeguards will slow innovation and deny benefits to people who need them. Some safeguards do impose costs, and poorly designed rules can block useful applications. Yet the absence of safeguards also has costs: errors, discrimination, privacy violations, loss of trust, and expensive remediation. A proportionate approach distinguishes low-risk assistance from high-impact decisions, supports testing and learning, and involves affected communities. Responsible innovation is not innovation without constraints; it is the capacity to create durable public benefit while taking foreseeable harms seriously.

22. Recommendations for Christian Leaders and Educators

Christian leaders should begin with theological literacy about technology. They need not become engineers, but should understand the basic difference between a model's output and human understanding, between a prediction and a moral judgment, and between technical performance and social value. Seminaries and church leadership programs can include AI ethics, privacy, misinformation, labor, and digital pastoral care. Case-based teaching is particularly useful: learners can examine a hypothetical church chatbot, an AI hiring tool, or a synthetic political video and identify affected people, duties, risks, and alternatives.

Churches should adopt practical governance. A designated leader or committee can maintain an AI-use register, review high-impact applications, and ensure staff training. Confidential pastoral data should not be entered into public systems without appropriate safeguards. Generated theological or public-facing material should be reviewed by a responsible human. Congregations should disclose synthetic media and avoid impersonation. Where AI is used in counseling or spiritual support, boundaries and escalation routes should be explicit. The church's commitment should be to facilitate human care, not to create a cheaper substitute for it.

Educators should teach students to use AI critically and honestly. Policies should state what assistance is allowed, how it should be acknowledged, and what counts as independent work. Rather than relying only on automated detection, educators can ask students to explain decisions, document revisions, compare sources, and connect knowledge to lived or professional contexts. Institutions should protect student data and avoid vendor contracts that compromise privacy or create unequal access. Faculty should receive time and support to redesign assignments rather than being expected to solve every issue individually.

Christian leaders should also engage in public policy and interfaith collaboration. AI affects people across religious boundaries; justice requires broad participation. Churches can work with Muslim, Hindu, Buddhist, and secular organizations, professional associations, universities, and community groups to promote digital literacy, accessibility, labor dignity, and responsible public systems. Public witness is strongest when it is grounded in service and careful evidence. The aim is to help society build technologies that enable people to flourish, especially those whose voices are often overlooked.

23. Conclusion: Technology under the Measure of Human Dignity

Artificial intelligence is a powerful expression of human creativity, and it can help address important needs. It can extend access to knowledge, improve coordination, support professional work, and assist communities facing complex challenges. Yet capability does not settle the question of purpose. Systems can reproduce injustice, weaken privacy, obscure accountability, displace workers, distort truth, and create dependence. The decisive question is whether AI is governed and used in ways that honor persons and serve the common good.

Christian faith offers a durable foundation for answering that question. The biblical image of God affirms a dignity that no performance score can grant

and no machine can remove. Neighbor-love requires attention to those who bear the risks of technology. Justice challenges structures that concentrate benefits and hide burdens. Stewardship places limits on power and includes care for creation. Truthfulness demands verification and resistance to deception. Human responsibility insists that institutions remain answerable for the systems they choose.

These commitments do not require Christians to reject AI. They call for discernment, competent governance, and courageous limits. In practice, a dignity-centered approach states a legitimate purpose, includes affected communities, tests for unequal harm, protects privacy, preserves meaningful appeal, assigns responsibility, and reviews outcomes. It uses AI where it strengthens human capacities and declines it where the cost to dignity, justice, or trust is too high. It recognizes that the moral quality of a system depends not only on its architecture but on the institutions and relationships in which it operates.

The promise of AI should be measured not by how closely machines imitate human expression but by whether human beings become more capable of knowing truth, caring for neighbors, pursuing justice, doing meaningful work, and protecting creation. Christian communities can help ensure that technological advancement remains a servant of life rather than a substitute for human responsibility. In this task, faith does not compete with sound science. It asks science and society to remember for whom knowledge is pursued and to whom power must answer.

This discernment is ongoing because technologies, institutions, and social expectations change. A responsible decision today may need revision when a model is updated, when a new population is affected, or when evidence reveals an unanticipated harm. Christian institutions can cultivate this attentiveness through regular review, open discussion, and the willingness to listen to people whose experiences differ from those of decision-makers. They

can also teach young people that competence includes knowing when not to automate, when to ask for help, and when a human relationship matters more than a rapid answer. The work of discernment belongs to communities as well as individuals: it requires shared language, trustworthy procedures, and leaders prepared to explain their choices. Such practices embody humility, because they acknowledge that no institution sees every consequence in advance. They embody hope, because they assume that people can learn, correct failures, and redirect technical power toward service. A Christian account of technological progress is therefore measured not only by novelty or adoption, but by the quality of the relationships and institutions that emerge around it.

Glossary

Artificial intelligence (AI): Computational systems designed to perform tasks involving prediction, classification, language, perception, planning, or related capabilities.

Algorithm: A set of computational procedures or rules used to transform inputs into outputs.

Automation bias: The tendency to over-rely on an automated recommendation, even when contrary evidence or uncertainty is present.

Data minimization: The practice of collecting and retaining only the personal data necessary for a defined, legitimate purpose.

Dignity: The inherent worth of every person, understood in Christian theology as grounded in God's creative and loving relationship with humanity.

***Imago Dei*:** Latin for "image of God," a biblical and theological account of human identity and vocation rooted especially in Genesis 1:26–27.

Human agency: A person's capacity to understand, deliberate, choose, act, and answer for conduct within relationships and institutions.

Human oversight: Meaningful human authority to understand, monitor, challenge, modify, or stop an AI-supported process.

Generative AI: AI systems that generate text, images, audio, code, or other outputs based on patterns learned from training data.

Model: A computational representation learned or specified to produce predictions, classifications, or generated outputs.

Responsible AI: The design, deployment, and governance of AI in ways that respect rights, fairness, safety, privacy, accountability, and social and environmental well-being.

Sociotechnical system: The combined arrangement of technology, data, people, organizational procedures, incentives, and institutions through which an AI system operates.

Stewardship: The accountable care of resources and creation entrusted to human beings, rather than unrestricted ownership or exploitation.

Transparency: Provision of relevant, understandable information about an AI system's purpose, data, limits, and role in decisions.

References

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. https://doi.org/10.1145/3442188.3445922

Coeckelbergh, M. (2020). *AI ethics*. MIT Press. https://doi.org/10.7551/mitpress/12549.001.0001

Dicastery for the Doctrine of the Faith, & Dicastery for Culture and Education. (2025, January 28). *Antiqua et nova: Note on the relationship between artificial intelligence and human intelligence*. The Holy See. https://www.vatican.va/roman_curia/congregations/cfaith/documents/rc_ddf_doc_20250128_antiqua-et-nova_en.html

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines, 28*, 689–707. https://doi.org/10.1007/s11023-018-9482-5

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence, 1*, 389–399. https://doi.org/10.1038/s42256-019-0088-2

National Council of Churches. (2021). *New Revised Standard Version Updated Edition Bible*. Friendship Press. (Scripture citations in this article follow the NRSV Updated Edition.)

UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. https://unesdoc.unesco.org/ark:/48223/pf0000381137

Biblical Passages Referenced

Genesis 1:26–28; 2:7, 15; 11:1–9; Exodus 20:3–6, 8–11; Leviticus 19:18; Deuteronomy 6:6–7; 10:17–19; Proverbs 4:7; 11:13; Amos 5:11–12; Micah 6:8; Matthew 6:6; 20:1–16; 22:37–40; Mark 1:40–45; 12:29–31; Luke 8:43–48; 10:25–37; John 1:14; 8:32; 13:34–35; 14:6; Acts 5:29; Romans 8:19–23; 14:12; Galatians 5:13; Ephesians 4:25; James 1:5; 2:1–9.