




Rudy C Tarumingkeng : *Ethical AI and the Christian Imagination -
Designing with Compassion and Justice*

By:

[Prof Ir Rudy C Tarumingkeng, PhD](#)

Professor of Management NUP: 9903252922

Rector, Cenderawasih State University, Papua (1978-1988, dan
Head, AGRO Manokwari Campus, (now Universitas Papua Manokwari)

Coordinator, CIDA/DIKTI - SFU Burnaby BC Canada 1988-1991

Rector, Krida Wacana Christian University, Jakarta (1991-2000)

Chairman, Board of Professors, IPB-University, Bogor (2005-2006)

AI - Data Analyst, dan Chairman Academic Senate, IBM-ASMI, Jakarta 2024-

© RudyCT Academic Series

rudyclt75@gmail.com

31 March 2026

ETHICAL AI AND THE CHRISTIAN IMAGINATION: DESIGNING WITH COMPASSION AND JUSTICE

Theology of Co-Creation: Man and Machine in God's Creative Plan 1. Introduction: Why Ethical AI Needs the Christian Imagination

Abstract

This essay explores the theological meaning of ethical AI through the lens of the Christian imagination, arguing that artificial intelligence should be designed not merely for efficiency, scale, and optimization, but for compassion, justice, truthfulness, and human dignity. Contemporary AI governance frameworks increasingly emphasize human rights, fairness, transparency, accountability, and human oversight. UNESCO's Recommendation on the Ethics of Artificial Intelligence places human rights and dignity at the center of AI ethics; the OECD's updated AI Principles frame trustworthy AI in relation to human rights and democratic values; NIST's AI Risk Management Framework stresses trustworthiness and social responsibility; the World Health Organization requires ethics and human rights to remain central to AI in health; and the European Union's AI Act has begun phased implementation of a risk-

based regulatory regime. These developments show that the ethical challenge of AI is not peripheral but civilizational. ([UNESCO](#))

From a Christian perspective, however, regulation and compliance, though necessary, are not enough. The deeper issue is the moral imagination that guides design itself. The essay argues that the Christian imagination offers a distinctive way of seeing the human person and the technological world: every person bears the image of God, every neighbor is owed love and justice, truth is a moral good, wisdom is more than calculation, and power must be reordered toward service. Relevant biblical texts—including Genesis 1–2, Exodus 35, Micah 6:8, Luke 10, Matthew 25, John 13, and James 3—suggest that technology should be governed not by mastery alone but by stewardship, mercy, and righteousness.

The essay develops this argument in five movements. First, it clarifies the meaning of the Christian imagination as a form of moral vision shaped by Scripture. Second, it examines biblical foundations for compassion and justice in relation to technical making. Third, it considers how contemporary AI systems create new ethical pressures around bias, labor, truth, care, exclusion, and vulnerable communities. Fourth, it proposes principles for designing AI with compassion and justice, including participatory design, bias mitigation, explainability, accountability, accessibility, and the preservation of human responsibility. Fifth, it reflects on the temptations of technocracy and idolatry, arguing that Christian ethics must keep AI within the bounds of creaturely service.

The essay concludes that ethical AI requires more than well-written policies. It requires communities capable of imagining technology as a form of neighbor-love. AI becomes ethically credible not when it appears most autonomous, but when it most faithfully serves human dignity, protects the weak, supports truthful communication, and contributes to the common good. In Christian terms, designing AI with compassion and

justice is an act of moral discipleship within the wider task of loving God and neighbor in a digital age.

Keywords

Ethical AI; Christian imagination; compassion; justice; human dignity; imago Dei; stewardship; AI governance; neighbor-love; technological ethics.

Ethical AI and the Christian Imagination: Designing with Compassion and Justice

Artificial intelligence is no longer a marginal research topic or a specialized industrial tool. It is becoming part of the ordinary architecture of modern life: healthcare, education, labor management, communications, public administration, finance, policing, logistics, and media are all being reshaped by AI-enabled systems. That is one reason major international institutions have moved rapidly to develop ethical and governance frameworks. UNESCO's Recommendation on the Ethics of Artificial Intelligence describes AI governance as one of the most consequential challenges of the present time and grounds its framework in human rights, dignity, transparency, fairness, and human oversight. The OECD's AI Principles, adopted in 2019 and updated in 2024, likewise promote innovative and trustworthy AI that respects human rights and democratic values. NIST's AI Risk Management Framework is designed to help organizations manage risks to individuals, organizations, and society, while WHO's AI guidance for health insists that ethics and human rights must remain at the heart of design, deployment, and use. The European Union's AI Act, which entered into force on 1 August 2024, has already begun phased application of prohibited practices rules, AI literacy obligations, governance obligations for general-purpose AI, and later rules for high-risk systems. ([UNESCO](#))

These developments are important because they show that AI is not merely a technical matter. It is an ethical field of action. Yet even the best governance frameworks cannot by themselves answer a more fundamental question: what kind of moral imagination should guide the design of AI in the first place? Regulation can prohibit certain harms and require certain procedures, but it cannot on its own teach a society what a human person is, why the vulnerable matter, why truth should be protected, or why efficiency should sometimes yield to mercy. Those are questions of moral anthropology and spiritual vision.

That is where the Christian imagination becomes relevant. By “Christian imagination” I do not mean fantasy, escapism, or aesthetic religious symbolism. I mean a way of seeing reality shaped by Scripture, prayer, worship, and theological reflection. The Christian imagination learns to see the world as creation, not raw material; human beings as image-bearers, not merely data points; neighbors as persons to be loved, not users to be managed; institutions as accountable before God, not only before markets; and technology as a tool for stewardship and service, not an autonomous measure of the good. This imagination does not reject technical innovation. Rather, it places innovation within a larger moral horizon.

Such an imagination is badly needed in an age when design is often driven by speed, competition, extraction, engagement metrics, and market dominance. AI systems are shaped not only by algorithms but by the values embedded in their objectives, datasets, interfaces, incentives, and deployment contexts. If those values are thin, then even impressive systems can become ethically corrosive. A model can be accurate and still unjust; efficient and still inhumane; engaging and still manipulative. The central problem is therefore not only whether AI can do more, but whether societies can see more deeply.

The Christian tradition has resources for such depth. It teaches that power must be measured by service, not domination; that truth matters

because persons matter; that justice is inseparable from care for the poor, the stranger, and the excluded; and that wisdom is more than computational capacity. Pope Francis has applied this insight directly to AI. In his 2024 message "Artificial Intelligence and the Wisdom of the Heart," he warned that this historical moment risks becoming "rich in technology and poor in humanity," insisted that reflection must begin from the human heart, and argued that machines can process data, but human beings alone can make sense of it in a truly humane way. He also cautioned that algorithms are not neutral and that preventive ethical regulation is needed to forestall harmful, discriminatory, and socially unjust effects. ([Vatican](#))

The same basic concern appears in the Vatican's 2025 note *Antiqua et nova*, which states that the Church encourages science, technology, and the arts as part of human collaboration in perfecting visible creation, but also insists that human intelligence must not be collapsed into AI and that functionalist views of the human person are deeply inadequate. The note treats AI not merely as a productivity tool, but as an anthropological challenge because it can reshape how people understand intelligence, dignity, work, relationships, and truth. ([Vatican](#))

This essay therefore argues that ethical AI requires the Christian imagination because only a morally thick vision of the human person can keep technical design from becoming spiritually and socially thin. Designing with compassion and justice is not an optional layer added after engineering. It is a discipline of seeing the neighbor rightly before designing systems that will affect the neighbor's life.

2. The Christian Imagination as Moral Vision

The Christian imagination begins with creation. Genesis does not present the world as an accidental material field awaiting capture by the strongest intelligence. It presents the world as a good creation spoken into being by God. Light, land, waters, vegetation, animals, and humanity are each granted place and goodness before they become useful for

human goals. This means that the moral starting point for any Christian account of technology is gratitude, not entitlement. The world is first received, then cultivated. It is never merely available.

The Christian imagination also begins with the human being as image-bearer. Genesis 1:26–27 declares that humankind is made in the image and likeness of God. This truth grounds dignity prior to utility. A person is not valuable because he is economically efficient, digitally legible, or cognitively superior. He is valuable because he is created and called by God. That claim becomes especially important in an AI-shaped society because machine systems tend to classify, rank, and optimize. Such practices can be useful, but they become morally dangerous when they start to define what a person is worth.

The Christian imagination, then, resists reduction. It resists reducing the person to data, intelligence to output, communication to content generation, and community to connectivity. It insists that persons are bodily, relational, moral, and spiritual beings. They have histories, vulnerabilities, responsibilities, and callings. They are capable not only of problem-solving but of love, repentance, worship, forgiveness, patience, and hope. AI can imitate certain expressions of intelligence. It cannot replace the fullness of personal life.

This imaginative stance is not anti-technical. Scripture affirms craft, skill, and making. Exodus 35:31–35 speaks of Bezalel as filled with the Spirit of God, with skill, intelligence, and knowledge in every craft, for the work of building the tabernacle. Human making is not outside the life of faith. Yet the example is revealing: technical skill is holy when it serves worship, community, and the purposes of God. The problem, therefore, is never simply that people make tools. The problem is what kind of world those tools are shaping and what kind of people their makers are becoming.

In modern design cultures, imagination is often colonized by certain defaults: scale is assumed to be good, speed is assumed to be progress, friction is assumed to be failure, personalization is assumed to be care,

and data extraction is assumed to be normal. The Christian imagination interrupts those defaults. It asks whether speed may sometimes destroy attentiveness; whether friction may sometimes protect dignity; whether personalization may slide into manipulation; whether optimization may invisibilize the weak; and whether scale may magnify injustice.

To say that AI should be designed with compassion and justice is therefore to say that its makers need an imagination capable of seeing beyond immediate performance metrics. They need a way of seeing children not only as end-users but as vulnerable persons; workers not only as cost centers but as bearers of vocation; patients not only as cases but as neighbors in need of care; citizens not only as data subjects but as participants in the common good. That way of seeing is not produced by code alone. It is formed by moral traditions, communities, and disciplines of attention.

Christian imagination is one such formation. It does not provide ready-made code libraries. It provides the moral grammar by which designers, institutions, and societies can ask what faithful making looks like in relation to God and neighbor. Without such grammar, ethical AI risks becoming a checklist rather than a vision.

3. Biblical Foundations: Compassion, Justice, Truth, and Wisdom

If Christian imagination is to guide AI design, it must be anchored in specific biblical themes rather than vague goodwill. Four themes are especially important: compassion, justice, truth, and wisdom.

Compassion in Scripture is not mere sentiment. It is active regard for the vulnerable. The ministry of Jesus is saturated with compassion: he sees crowds “like sheep without a shepherd” (Mark 6:34), touches the sick, feeds the hungry, and refuses to let suffering remain abstract. In Luke 10, the Good Samaritan does not optimize charity from afar; he interrupts his own journey, bears cost, and takes responsibility for the wounded stranger. In Matthew 25, the moral test of nations is framed in terms of

care for the hungry, the thirsty, the stranger, the naked, the sick, and the prisoner. Christian compassion, then, is not a private emotion. It is a way of structuring response around the vulnerable.

Justice is equally central. Micah 6:8 condenses covenantal ethics into three demands: do justice, love mercy, walk humbly with God. Isaiah 1 condemns religious performance that coexists with oppression. Amos 5 refuses liturgical self-confidence where justice is absent. Proverbs repeatedly defends truthful judgment and warns against the exploitation of the poor. In biblical perspective, justice is not optional to faith. It is one of the primary ways faith appears publicly.

These themes matter directly for AI. Compassion asks what design choices do to people who are already exposed to risk. Justice asks how benefits and harms are distributed, who is excluded, whose data trains the system, who bears the cost of errors, and whether the voices of the poor and powerless are heard. Pope Francis highlighted this exact concern in a 2025 message on AI, recalling that debates about AI regulation must include “the poor, the powerless and others who often go unheard in global decision-making processes.” ([Vatican Press](#))

Truth is another central biblical concern. The commandment against false witness is not merely about courtroom perjury; it protects the social conditions of trust. John 8:32 links truth with freedom. Ephesians 4:25 commands truth-speaking because members belong to one another. AI design becomes morally serious here because generative systems can produce falsehood at scale, plausible fabrications, synthetic impersonation, and environments in which reality is distorted by engagement incentives. Pope Francis warned that deepfakes and false narratives can create “cognitive pollution,” distort reality, and corrode human relations. He argued that the technology of simulation becomes perverse when it distorts our relationship with others and with reality. ([Vatican](#))

Finally, wisdom is indispensable. Biblical wisdom is not merely intelligence. Proverbs treats wisdom as moral discernment grounded in reverence for God. 1 Kings 3 presents Solomon asking not for raw power but for “an understanding mind” to discern between good and evil. James 3 distinguishes wisdom from selfish ambition. Wisdom is therefore integrative: it judges not only what can be done, but what should be done and toward what end.

This distinction has become central in recent Christian reflection on AI. Pope Francis’s “Wisdom of the Heart” message says explicitly that wisdom cannot be sought from machines, that the term “artificial intelligence” can be misleading, and that human beings alone can make sense of data in a truly human way. He adds that everything within human reach becomes either an opportunity or a threat depending on the inclination of the heart. ([Vatican](#))

Taken together, these themes offer a strong biblical grammar for ethical AI. Compassion directs design toward the suffering and excluded. Justice requires fairness, accountability, and participation. Truth resists deception and epistemic manipulation. Wisdom keeps technical power under moral judgment. This is not an external Christian ornament added to secular design. It is a coherent moral vision for making systems that affect human life at scale.

4. Why AI Intensifies the Demands of Compassion and Justice

Every era has had its tools, but AI intensifies ethical demands because it combines scale, speed, opacity, and institutional reach. A biased human decision can be harmful. A biased model embedded across hiring, lending, healthcare, education, and public administration can multiply harm across millions of people. A misleading message can deceive a few; generative systems can flood information environments with plausible falsehoods at industrial scale. A poorly designed interface can frustrate users; a large AI platform can reshape attention, learning, trust, and labor markets.

UNESCO explicitly frames this problem in such terms. Its AI ethics page says AI creates many opportunities, but also raises profound ethical concerns because AI systems can embed bias, contribute to climate degradation, threaten human rights, and compound existing inequalities, thereby causing further harm to already marginalized groups. It identifies human rights and human dignity as the cornerstone of the Recommendation and includes policy action areas spanning data governance, environment and ecosystems, gender, education, research, and health and social wellbeing. ([UNESCO](#))

The OECD's updated Principles similarly show that AI ethics is broader than narrow technical reliability. They emphasize innovative and trustworthy AI that respects human rights and democratic values. This indicates that AI must be judged not only by whether it works, but by whether it works in a way compatible with freedom, fairness, pluralism, and the dignity of persons and communities. ([OECD](#))

NIST's AI RMF makes a parallel point in a more operational register. It describes the framework as a resource for managing risks to individuals, organizations, and society associated with AI, and states that risk management should help cultivate trustworthiness and public trust. This matters because trust is not generated merely by technical sophistication. Trust depends on whether people believe systems are safe, fair, accountable, and designed with regard for their wellbeing. ([NIST](#))

WHO's guidance for health shows how concrete these concerns become in high-stakes settings. The guidance recognizes AI's promise to improve diagnosis, treatment, research, and public health functions, but insists that ethics and human rights must be central to design, deployment, and use. It identifies six consensus principles and recommends governance measures that keep stakeholders accountable and responsive to healthcare workers, communities, and individuals affected by AI. ([World Health Organization](#))

The EU AI Act likewise illustrates the political seriousness of these issues. Its phased implementation—beginning with prohibited AI practices and AI literacy obligations in February 2025, governance rules and obligations for general-purpose AI models in August 2025, and later obligations for certain high-risk systems—shows that policymakers increasingly treat AI not as a voluntary ethics matter alone but as a field demanding enforceable public safeguards. ([Digital Strategy](#))

For Christian ethics, the implication is clear: compassion and justice are not softer alternatives to “serious” design. They are exactly what serious design requires in a world where technical decisions can become structural moral realities. If AI can magnify power, then compassion and justice must be designed upstream rather than added downstream as public relations language.

5. Designing with Compassion

What would it mean, concretely, to design AI with compassion? Compassion in design begins with attentive seeing. In the Gospels, Jesus repeatedly notices those whom systems overlook: the sick, the shamed, the poor, the outsider. Design shaped by compassion therefore begins not from the ideal user, but from the vulnerable user. It asks who will be harmed first if the system fails, who may be excluded by default assumptions, who lacks the resources to contest errors, and who is likely to be invisibilized by efficiency logics.

A compassionate design process therefore requires listening before deployment. A system built for education should involve not only engineers and administrators but teachers, students, parents, and especially those in under-resourced contexts. A system built for healthcare should involve patients, nurses, clinicians, caregivers, and communities likely to bear the greatest risk. Participatory design is not merely a democratic preference; it is a form of moral attention. It operationalizes the truth that affected people are not afterthoughts.

UNICEF's updated 2025 work on AI and children is especially instructive here. UNICEF warns that children face urgent risks ranging from AI-generated disinformation to harmful deepfakes and exploitative content, while remaining largely excluded from shaping the systems that affect them, particularly in low-resource settings. It says evidence is still limited on how AI affects children's social, emotional, and cognitive development and emphasizes that child-centred AI policies and systems must place children's rights at the heart of the AI age. ([UNICEF](#))

A Christian account of compassion takes such concerns seriously because children occupy a privileged moral place in the Gospels. Jesus warns against causing little ones to stumble and treats care for the vulnerable as a sign of the kingdom. Designing AI with compassion, then, requires more than age gating or parental settings. It requires a fundamental bias toward protection, developmental appropriateness, and non-exploitation where children are concerned.

Compassion also requires preserving human presence in contexts of suffering. WHO's guidance and the Vatican's AI note both point in this direction. WHO says AI for health must remain accountable to healthcare workers and affected communities. *Antiqua et nova* affirms that AI can support diagnosis and treatment, but warns that if the patient-provider relationship is replaced rather than enhanced, an essential human structure can be reduced to something fragmented and impersonal. ([World Health Organization](#))

That point is theologically crucial. Compassion is not just correct output. It is relational attention. A chatbot may provide immediate information, but it does not thereby fulfill the moral reality of care. A compassionate AI design ethic will therefore distinguish between assistance and substitution. It will welcome tools that reduce clerical burdens, extend access, and support decision-making, while resisting the temptation to let institutions use automation as an excuse to withdraw embodied care from those who most need it.

Compassion in design also means attending to failure. Systems fail. Models hallucinate. Interfaces mislead. Ranking systems misclassify. A compassionate system is one designed not only for performance when all goes well, but for recoverability when things go wrong. It provides appeal paths, human contact, and graceful fallback mechanisms. It does not trap already-vulnerable users inside impersonal error loops.

Finally, compassion requires slowness at certain points. Not every decision should be accelerated. In pastoral, medical, legal, or educational settings, speed may degrade prudence. The Christian imagination recognizes that neighbor-love sometimes means delaying automation in order to preserve responsibility and attentiveness. In this sense, compassionate design is not anti-innovation. It is anti-cruelty.

6. Designing with Justice

If compassion directs attention to the vulnerable, justice directs structure toward fairness and right relation. Designing AI with justice means asking how power, benefit, and burden are distributed throughout the AI lifecycle: data collection, labeling, modeling, deployment, monitoring, and redress.

The first justice question concerns data. Whose behavior is represented? Who is missing? Who was over-surveilled in order to produce the dataset? Whose language, accent, culture, disability, or social context is treated as marginal noise? Justice requires more than balancing classes in a spreadsheet. It requires historical awareness. Many datasets inherit older inequities, institutional prejudices, and colonial asymmetries. To train on such histories without reflection is to automate memory without repentance.

UNESCO's emphasis on fairness, inclusiveness, and human rights, alongside action areas such as data governance and gender, reflects this structural concern. Its framework does not stop at abstract principle; it

recognizes that ethical AI requires institutional attention to the concrete places where injustice enters systems. ([UNESCO](#))

The second justice question concerns explainability and contestability. A Christian ethic of justice cannot accept systems that impose serious consequences on persons without meaningful avenues for understanding and response. Proverbs praises just weights and measures; biblical justice is not comfortable with opaque mechanisms that determine outcomes while concealing criteria. This does not mean every model must be fully intelligible to every user in a technical sense. It does mean institutions must preserve accountability, explanations suitable to the stakes involved, and channels for challenge and remedy.

Here the NIST AI RMF is useful because it ties trustworthiness to context, governance, and risk management rather than assuming one generic standard for all systems. The EU AI Act's risk-based structure points in the same direction by differentiating obligations and treating some uses as especially sensitive or unacceptable. ([NIST](#))

The third justice question concerns labor. AI systems are often praised for productivity, but productivity can conceal exploitation. Workers may be surveilled more closely, deskilled, pushed to algorithmic pace, or displaced from meaningful roles without voice in the redesign of work. Christian ethics cannot ignore this because work is not just income generation. It belongs to vocation, social participation, and dignity. A just AI design ethic asks whether systems augment workers, protect them, retrain them, and share gains fairly, or whether they merely extract more value from labor while shifting risk downward.

This concern is also present in Christian institutional reflection. The Rome Call for AI Ethics urges that AI be developed for the good of humanity rather than the gradual replacement of human beings, and it insists that technological progress should affirm human brilliance while remaining dependent on ethical integrity. It also places inclusion, dignity, and environmental care at the heart of the ethical outlook.

The fourth justice question concerns voice. Pope Francis's warning that the poor and powerless must be included in AI debates is especially relevant here. Justice in design is not only about measuring harms after deployment; it is about who is allowed to define the problem before deployment. If AI systems are designed primarily by wealthy actors for wealthy environments, then global inequalities will be reproduced even when the systems seem technically successful. ([Vatican Press](#))

Finally, justice in AI must include ecological justice. UNESCO includes environment and ecosystem flourishing among the four core values of its Recommendation, and the Rome Call similarly insists that AI should serve humanity and the environment, our "common and shared home." This matters because digital systems are not immaterial in their effects. They depend on infrastructure, energy, supply chains, and waste streams. Designing with justice therefore includes asking who bears the ecological cost of digital abundance. ([UNESCO](#))

7. Truthful Design and the Protection of Reality

Compassion and justice both depend on a truthful environment. A society cannot care well for the vulnerable or distribute goods fairly if its informational life is poisoned by fabricated media, plausible falsehoods, impersonation, or manipulative recommender systems. That is why truthful design must be counted as part of ethical AI.

Pope Francis's 2024 communications message is especially powerful on this point. He notes that AI can aid communication and make knowledge more accessible, but also warns of "cognitive pollution," fake news, deepfakes, and simulations that distort relations with others and with reality. He argues that the digital revolution can widen freedom only if it does not imprison people in echo chambers and unverified narratives. He explicitly asks how societies can make the criteria guiding algorithms more transparent, how they can identify the authorship of texts and sources, and how they can distinguish whether an image or video portrays an event or simulates it. ([Vatican](#))

This concern converges with emerging regulatory efforts. The EU's AI implementation materials note transparency obligations related to AI-generated content, including marking generated content and disclosing the artificial nature of images, audio, and text. Such measures are not merely administrative. They reflect the moral importance of distinguishing reality from simulation. ([Digital Strategy](#))

Christian theology can explain why this matters so deeply. Truth in Scripture is not only factual correctness. It is covenantal faithfulness. False witness breaks community because it weaponizes speech against the neighbor. If AI systems make deception cheap, scalable, and difficult to detect, then they threaten not only epistemic order but moral order. Designing AI ethically therefore requires more than content safety filters. It requires a design stance committed to traceability, disclosure, restraint in synthetic impersonation, and humility about the effects of simulation.

Truthful design also resists anthropomorphic deception. When systems are built to appear more human than they are, users may be nudged into misplaced trust, dependency, or emotional attachment under false assumptions. Theologically, deception about agency matters because it blurs responsibility. A Christian design ethic should therefore prefer honesty about the nature and limits of systems over the theatrical illusion of personhood.

Moreover, truthful design must protect plurality. Pope Francis warns against groupthink and the collapse of pluralism into data-driven uniformity. Ethical AI should therefore be designed not simply to maximize engagement or reinforce user certainty, but to support a more honest encounter with reality's complexity. ([Vatican](#))

In a machine-mediated society, truth is itself a justice issue. Those with fewer resources are often less able to defend themselves against reputational harm, algorithmic misrepresentation, or informational manipulation. Designing for truth is therefore part of designing for compassion and justice.

8. Human Oversight, Accountability, and the Limits of Automation

One of the most important themes across contemporary AI ethics frameworks is human oversight. UNESCO identifies it as a core principle. NIST's framework is built around governance and accountable risk management. WHO insists that stakeholders remain accountable and responsive. The OECD links trustworthy AI to democratic values and human rights. The EU AI Act operationalizes some of these concerns through duties, governance structures, and risk categories. ([UNESCO](#))

Christian ethics not only agrees with this emphasis; it deepens it. Human oversight matters because moral agency belongs to persons, not tools. In the biblical world, judgment, repentance, mercy, and accountability are not outsourced to mechanisms. Even where institutions and procedures are necessary, responsibility remains personal and communal. One cannot say, in any morally satisfying way, "the algorithm decided," and thereby end the ethical conversation.

This is one reason Pope Francis, at the G7 in 2024, insisted that AI is above all a tool and that its benefits or harms depend on its use. He warned that its growing influence on medicine, work, communications, education, and politics makes serious reflection unavoidable, and he stressed that proper human control over AI-mediated choices must be safeguarded because human dignity is at stake. He also warned of the danger that AI might deepen injustice between advanced and developing nations or between dominant and oppressed classes if not governed wisely. ([Vatican](#))

For Christian imagination, then, oversight is not just a technical governance box. It is a moral refusal to let the appearance of intelligence eclipse the reality of responsibility. Designers remain responsible for objectives and interfaces. Institutions remain responsible for deployment contexts and redress mechanisms. Leaders remain responsible for whether systems are being used to serve or to dominate. Citizens remain

responsible for whether they accept convenience at the cost of the vulnerable.

At the same time, “human oversight” must not become a hollow phrase. Superficial oversight—where a human signs off on outputs they cannot really question, or where responsibility is nominally assigned but structurally diffused—is ethically weak. Genuine oversight requires competence, authority, time, and organizational support. It also requires the courage to slow down or withdraw a system when harms emerge.

A Christian ethic of accountability would add one more layer: confession and correction. Institutions should be capable of admitting that a system was unjust, misleading, or harmful. Without that moral capacity, oversight becomes defensive reputation management. Because humans are fallible, ethical design must include not only governance before harm but practices of truth-telling after harm.

9. Children, the Poor, and the Margins as Design Priority

The strongest test of Christian AI ethics is not how well it serves the powerful, but how faithfully it protects those who are easiest to ignore. Throughout Scripture, God’s judgment falls heavily on societies that neglect the poor, exploit workers, silence the weak, or turn justice into an advantage for the already secure. Any Christian account of AI that does not take this priority seriously has not yet become morally serious.

UNICEF’s 2025 child-centered AI work is instructive again here. It notes that children are both heavy users of AI-shaped systems and still largely excluded from influencing the systems that affect them. It highlights risks such as exploitative content, deepfakes, disinformation, and developmental uncertainty, while calling for child rights to be central in AI policy and design. ([UNICEF](#))

That principle extends beyond children. The poor, migrants, people with disabilities, linguistic minorities, rural communities, the digitally excluded, and those in low-resource settings often encounter AI systems in the

most coercive forms: welfare assessment, border control, predictive policing, automated hiring filters, low-wage platform labor, and administrative gatekeeping. A system may appear neutral in affluent contexts and still be harsh in marginal ones.

Pope Francis's statement that the poor and powerless must be heard in AI regulation is especially important because it rejects a narrow model of expertise. Technical expertise matters, but ethical legitimacy also depends on whether those with the least power are visible in the design imagination. ([Vatican Press](#))

For Christian ethics, this priority is theological before it is political. In Matthew 25, Christ identifies himself with the hungry, the stranger, the naked, the sick, and the imprisoned. In other words, the face of the vulnerable becomes a place of encounter with Christ. Designing AI with compassion and justice therefore cannot mean abstract benevolence toward "users" in general. It means concrete regard for those whose interests are easiest to discount when systems are optimized for convenience, compliance, or profit.

In practical terms, this implies at least four design commitments. First, impact assessments should explicitly consider vulnerable populations rather than assume generic users. Second, affected communities should be represented in consultation and evaluation processes. Third, systems deployed in high-stakes contexts should include robust redress and human review for those least able to navigate bureaucratic systems. Fourth, metrics of success should include protection from harm, not just adoption, throughput, or cost savings.

Christian imagination here becomes a discipline of reversal. It asks designers to begin where Christ so often begins: with the one left by the roadside, the one not invited to the center, the one whose suffering is easiest to abstract away. AI designed from the margins will look different from AI designed only for powerful institutions.

10. The Temptations of Technocracy and Idolatry

Any constructive Christian ethic of AI must also diagnose spiritual temptations. Two are especially dangerous: technocracy and idolatry.

Technocracy is the habit of assuming that technical capability and managerial control are sufficient to define the good. Under technocracy, questions of purpose, meaning, mercy, and justice are translated into optimization problems. What matters becomes what can be measured. What cannot be measured becomes secondary. In such a worldview, compassion can be reduced to sentiment, justice to compliance, and wisdom to prediction.

Pope Francis has repeatedly criticized this tendency. In the 2024 “Wisdom of the Heart” message, he warns against reducing individuals to data, thinking to mechanical process, and goodness to profit. He insists that information cannot be separated from living relationships and that the representation of reality in big data, however useful for machines, involves a substantial loss of the truth of things if it cuts reality off from embodied human experience. ([Vatican](#))

The Vatican’s *Antiqua et nova* echoes the warning in a different register. It cautions against functionalism, against equating human intelligence too closely with AI, and against evaluating persons by what they can perform. It also notes that AI can only simulate authentic relationships and cannot substitute for the relational and moral development proper to the human person. ([Vatican](#))

Idolatry goes deeper still. It occurs when a created thing is treated as ultimate, when human beings begin to trust a tool in ways that displace reverence, moral judgment, and dependence on God. AI becomes idolatrous not only when people speak in exaggerated language about “superintelligence,” but whenever technical systems silently become the standard by which humanity judges itself. Then prayer looks

unproductive, contemplation inefficient, embodied care too slow, and moral struggle inferior to optimization.

The Christian imagination resists this by remembering creatureliness. Human beings are not gods. Machines are not saviors. No system, however capable, can answer the deepest questions of meaning, sin, forgiveness, holiness, or love. Ethical AI therefore requires spiritual realism. It requires communities capable of using powerful tools without surrendering their understanding of the good to those tools.

This does not imply fear of technology. It implies right order. A hammer is useful when it remains a hammer. A recommender system is useful when it remains accountable. A diagnostic model is useful when it remains a servant of care. When technologies begin to define what persons are for, the order has been reversed. Christian ethics calls that reversal by its proper name.

11. A Constructive Christian Framework for Ethical AI

A constructive Christian framework for AI design can be summarized in seven commitments.

First, **dignity before utility**. Every person affected by AI remains an image-bearer whose value precedes system goals. This principle should shape datasets, interfaces, deployment decisions, and redress mechanisms.

Second, **compassion before convenience**. Systems should be designed first around those most vulnerable to error, exclusion, or exploitation. The question is not only what is easiest to scale, but what is safest and most humane for those at risk.

Third, **justice before efficiency**. Gains in speed or cost reduction do not justify discriminatory outcomes, opaque decisions, labor degradation, or silence toward the structurally weak.

Fourth, **truth before engagement**. AI systems that communicate, generate, or rank content should be designed to protect traceability, disclosure, and reality-based understanding rather than merely maximizing consumption.

Fifth, **wisdom before automation**. Not every human judgment should be delegated. High-stakes contexts—medicine, law, pastoral care, education, public administration—require discernment about which tasks may be assisted and which should remain irreducibly personal.

Sixth, **community before control**. AI should strengthen human relationships, not displace them. It should support journalists rather than abolish on-the-ground reporting, support clinicians rather than replace care, support teachers rather than flatten education into content delivery, and support workers rather than turn them into adjuncts of machine pace.

Seventh, **stewardship before domination**. Technology belongs within a larger responsibility to creation, the common good, and future generations. That includes ecological cost, labor conditions, and the distribution of digital power.

This framework resonates with many public principles already articulated by UNESCO, OECD, NIST, WHO, the EU, and the Vatican. But its distinctively Christian contribution lies in its grounding. It does not defend dignity because dignity is rhetorically useful. It defends dignity because persons are made in God's image. It does not defend compassion because empathy is trendy. It defends compassion because Christ identifies with the vulnerable. It does not defend justice because unfair systems are inefficient. It defends justice because righteousness is part of faithful life before God.

In this sense, designing AI with compassion and justice is not merely ethical engineering. It is moral discipleship under technological conditions.

12. Conclusion

Ethical AI and the Christian imagination belong together because AI design is finally a question of how human beings see one another. If they see one another as profiles, optimization targets, labor inputs, or managed populations, then AI will likely be built in the image of domination. If they see one another as neighbors, image-bearers, and vulnerable persons whose dignity is inviolable, then AI can be designed more faithfully toward service.

Contemporary governance frameworks already point in this direction. UNESCO centers human rights, dignity, fairness, transparency, human oversight, and environmental flourishing. The OECD frames trustworthy AI in relation to democratic values and rights. NIST emphasizes trustworthiness and risk management for individuals and society. WHO keeps ethics and human rights at the center of AI in health. The EU AI Act translates some of these concerns into a staged legal regime. The Vatican's recent documents add a theological depth by warning against functionalism, technocracy, and the loss of humanity in a machine-driven culture. ([UNESCO](#))

Yet none of these frameworks can do the whole work alone. The decisive question remains whether designers, institutions, churches, schools, and publics possess an imagination broad enough to keep compassion and justice at the center. Christian ethics insists that they must. It teaches that power must serve the weak, that truth must be guarded for the sake of freedom, that labor must not be stripped of dignity, that children must not be sacrificed to novelty, and that technology must remain a creaturely instrument rather than an idol.

The future of ethical AI will not be decided only in laboratories or legislatures. It will also be decided in the moral habits by which communities learn to ask: Who is my neighbor here? Who is being harmed? Who has been left out? What does love require? What does

justice require? Those questions are ancient, but they have become newly urgent in a digital age.

If AI is designed with compassion and justice, it can become part of a more humane technological order—one in which technical brilliance is disciplined by mercy, fairness, truthfulness, and wisdom. In Christian terms, that would mean that AI had begun to serve not the illusion of omnipotence, but the real work of loving God and neighbor in history.

Glossary

AI ethics

The field of reflection and governance concerned with how artificial intelligence should be designed, used, regulated, and limited in relation to persons, institutions, and society.

Christian imagination

A morally formative way of seeing reality shaped by Scripture, worship, prayer, and theological reflection, through which human beings interpret persons, institutions, and technologies.

Compassion

Active regard for the suffering and vulnerability of others, expressed not only in feeling but in concrete forms of care, protection, and responsibility.

Data governance

The policies, structures, and practices by which data are collected, managed, used, protected, and made accountable.

Explainability

The degree to which an AI system's output or behavior can be understood, interpreted, and meaningfully questioned in context.

Functionalism

A view that defines intelligence or personhood primarily by what an entity does rather than by its deeper moral and ontological reality.

Human oversight

The principle that human beings remain responsible for supervising, directing, reviewing, and, when necessary, overruling AI systems.

Imago Dei

Latin for "image of God," referring to the biblical teaching that every human being bears God-given dignity and vocation.

Justice

Right relation in social, institutional, and personal life; in this essay, justice includes fairness, accountability, inclusion, redress, and special concern for the vulnerable.

Neighbor-love

The Christian moral command to love others concretely and responsibly, especially those whose need places a claim upon one's action.

Technocracy

A mindset in which technical efficiency, optimization, and managerial control are treated as sufficient guides to human flourishing.

Truthful design

A design orientation that protects reality-based understanding, disclosure, accountability, and resistance to deception or manipulative simulation.

Wisdom

Moral discernment that judges not only what can be done, but what should be done and toward what end; in Christian theology, wisdom is more than information processing.

References (APA 7)

- Dicastery for the Doctrine of the Faith, & Dicastery for Culture and Education. (2025). *Antiqua et nova: Note on the relationship between artificial intelligence and human intelligence*. Vatican. ([Vatican](#))
- European Commission. (2025). *AI Act: Shaping Europe's digital future*. European Commission. ([Digital Strategy](#))
- Francis. (2024a). *Artificial intelligence and the wisdom of the heart: Towards a fully human communication*. Vatican. ([Vatican](#))
- Francis. (2024b). *Address of His Holiness Pope Francis at the G7 session on artificial intelligence*. Vatican. ([Vatican](#))
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce. ([NIST](#))
- Organisation for Economic Co-operation and Development. (2024). *AI principles*. OECD. ([OECD](#))
- Pontifical Academy for Life. (2020). *Rome Call for AI Ethics*. Vatican.
- The Holy Bible, New Revised Standard Version Updated Edition. (2021). National Council of the Churches of Christ in the United States of America.
- United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO. ([UNESCO](#))
- United Nations Children's Fund. (2025). *Guidance on AI and children 3.0*. UNICEF Innocenti. ([UNICEF](#))

- World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. World Health Organization. ([World Health Organization](#))

Copilot for this article: ChatGPT 5.2. Thinking (2025). Access date: 31 March 2026. Author's account

<https://chatgpt.com/c/69cb206f-7500-839a-9b8c-584686bcd49d>