

MAGNIFICA HUMANITAS AND ONE HUMANITY

A Policy Paper on Human Dignity, Artificial Intelligence, Social Justice, and the Future of Humanity

*Prepared on the basis of the encyclical letter *Magnifica Humanitas* and the *One Humanity* books by Alexandre Caldas*

Alexandre Caldas, New York, 25 May 2026

Abstract

Magnifica Humanitas is significant because it places the word *Humanity* at the center of a new moral vocabulary for the AI era. Its central claim is not anti-technology. It is a demand that technological power be governed by a doctrine of the human person: dignity before efficiency, the common good before extraction, solidarity before domination, subsidiarity before centralization, justice before speed, and truth before manipulation. The letter interprets artificial intelligence as one of the *res novae* of our time: not simply a tool to be regulated after deployment, but a civilizational path that must be oriented toward the good from its earliest conception, financing, *design*, training, governance and use.

The *One Humanity* framework proposed in the two books converges strongly with this diagnosis. It is also a response to a civilizational moment in which geopolitics, climate, poverty, digital power, and institutional weakness interact as systemic risks. It argues that a future world order must be redesigned around people, places and planet; that the poorest nations and Small Island Developing States must move from the margins to the center of global priority; and that technology must be bound to dignity, accountability and human agency.

The policy proposition of this paper is therefore direct: *Magnifica Humanitas* provides the ethical and spiritual language for human-centered AI; *One Humanity* provides an operational architecture for global cooperation. Together, they can support a doctrine of technological social justice: artificial intelligence designed for human dignity, deployed for the common good, governed through subsidiarity, measured through human impact, and financed in solidarity with those most vulnerable to the disruptions of the AI age.

Core recommendation for Future Action:

Create a One Humanity AI and Social Justice Compact: a practical framework translating the principles of Catholic Social Doctrine and the One Humanity architecture into mandatory design, procurement, governance and impact-assessment standards for AI systems affecting rights, work, public services, education, peace, media and democratic life.

1. The encyclical as a courageous voice for the future

The courage of *Magnifica Humanitas* lies in its refusal to reduce AI ethics to a technical compliance problem. The letter frames AI (as one among emerging technologies) as a question about the destiny of the human person and the future of the human family. It begins with a civilizational choice: the *Tower of Babel* or the *Rebuilding of Jerusalem*. Babel stands for technological self-assertion, uniformity, pride, the sacrifice of dignity to efficiency and the temptation to translate the whole mystery of the person into data and performance. Jerusalem represents shared responsibility, prayerful discernment, plurality, communion, and the rebuilding of relationships before the rebuilding of systems.

This biblical contrast is a policy framework. It asks whether the AI era will become a world of centralized private power, opaque decision systems, automated exclusion, synthetic manipulation, and algorithmic control; or whether it can become a new architecture of social justice, where technology heals, connects, educates and protects the common home. The letter does not deny the benefits of technology. It insists that technology is never neutral in practice because it reflects the values, incentives, funding, institutions, and imaginaries of those who design and deploy it.

Its political significance is equally strong. The letter identifies a shift in power from the public authority of States toward private and transnational technological actors whose resources can surpass the capacities of many governments. This creates a governance asymmetry. The AI era is not only about algorithms; it is about who controls the infrastructure of intelligence, who defines the metrics of social value, who owns the data, who benefits from automation, and who bears the harms.

The encyclical is therefore courageous because it speaks to a new concentration of power with an older but still radical language: human dignity, the common good, universal destiny of goods, subsidiarity, solidarity, and social justice. These principles are not decorative. They are criteria for *institutional design*, *AI design*, industrial policy, public procurement, education, economic transition, and peacebuilding.

Babel versus Jerusalem as AI governance models

Dimension	Babel model	Jerusalem model
Power	AI as dominance, extraction, and uniformity	AI as service, care, and shared responsibility
Design	Efficiency-first optimization	Dignity-by-design and social impact-by-design
Governance	Private opacity and regulatory lag	Transparent accountability with public interest oversight
Human person	User, data point, worker to automate	Person, conscience, relationship, and vocation
Social outcome	Fragmentation, dependency, exclusion	Communion, inclusion, human flourishing

2. Humanity and human dignity as the doctrinal center

The repeated use of Humanity is the letter's theological and political center of gravity. Humanity is not treated as an abstraction. It is the concrete unity of persons whose dignity precedes markets, States, platforms, and machines. The human being is not valuable because of productivity, data generation, cognitive performance, or economic utility. Human dignity is intrinsic, equal, non-transferable and prior to every system.

This emphasis is decisive for the AI era. AI systems are often built around proxies: scores, probabilities, profiles, behavior predictions, risk classes, and optimization functions. These proxies can be useful, but they can also become reductions of the human person. A human-centered doctrine insists that no person may be fully captured by a model, and no automated system should be allowed to substitute probabilistic inference for moral judgment where fundamental rights, work, health, education, welfare, migration, policing or democratic participation are at stake.

The social doctrine principles as AI design principles

- **Human dignity:** AI must preserve agency, privacy, conscience, relational life, and the right not to be reduced to a data profile.
- **Common good:** AI must be evaluated by its contribution to social trust, public goods, inclusion, peace, education, health, and ecological sustainability.
- **Universal destination of goods:** data, knowledge, compute, and AI benefits cannot be monopolized in ways that deepen exclusion between rich and poor nations.
- **Subsidiarity:** decisions should be made at the closest competent level, with affected communities, workers, teachers, patients and citizens participating in design and oversight.
- **Solidarity:** AI transition costs must not be pushed onto the vulnerable; benefits and protections must be shared across societies and generations.
- **Social justice:** AI must be assessed through distributive, procedural and corrective justice, including bias remediation, access, redress, and accountability.

These principles transform AI governance from a narrow question of safety to a broader doctrine of justice. A system can be technically safe and still socially unjust. It can be accurate and still dehumanizing. It can be efficient and still violate subsidiarity by centralizing decisions away from persons and communities. The encyclical's contribution is to insist that the moral question is not only whether AI works, but what kind of humanity it builds.

3. AI concerns: power, truth, work, freedom, and peace

Magnifica Humanitas identifies AI as a valuable tool requiring vigilance. Its concern is not only malfunction, but misdirection. The deepest risk is a technocratic paradigm in which digital power becomes a substitute for moral wisdom and social judgment. Four domains require special attention.

Truth and communication

AI can strengthen knowledge, translation, scientific discovery, and access to information. It can also flood the public sphere with synthetic content, manipulate attention, intensify polarization, and erode the common good of truth. The result is not merely misinformation but an altered collective imagination, where citizens no longer share a reliable basis for democratic deliberation. A policy response must therefore move beyond content moderation toward an ecology of communication: provenance, media literacy, transparency of synthetic content, responsibility of platforms, and public-interest infrastructures for trustworthy knowledge.

Work and economic transition

AI will reshape work. The issue is not simply job displacement, but the meaning and dignity of work. Work is a site of contribution, social identity, creativity, responsibility, and solidarity. An AI economy that treats labor as a cost to eliminate weakens the moral foundations of society. A just AI transition requires worker participation, reskilling, social protection, new models of shared productivity gains, and investment in sectors where human care, education, creativity, ecological restoration, and community service are valued.

Freedom and dependency

AI systems can personalize services and improve accessibility, but they can also commercialize attention, predict vulnerabilities, automate behavioral nudges, and build dependencies. The protection of freedom requires limits on manipulative design, exploitative data extraction, addictive architectures, surveillance-based business models, and algorithmic forms of social control.

Peace, war, and multilateralism

The letter’s concern with weapons and artificial intelligence is central. Autonomous or semi-autonomous systems, targeting algorithms, cyber operations, and AI-enabled information warfare risk normalizing force without accountability. The crisis of multilateralism intensifies this risk. AI governance must therefore include arms-control frameworks, human responsibility for lethal decisions, confidence-building measures, verification systems, and renewed diplomacy. Peace is not only the absence of war; it is the institutionalization of justice before violence becomes rationalized.

4. Embedding principles into the design phase of AI

The strongest policy implication of *Magnifica Humanitas* is that ethical principles must be built into the design phase, not added as public relations after deployment. Human dignity must become an engineering, procurement, and governance requirement. This means moving from “ethics by declaration” to “justice by design.”

A human dignity-by-design lifecycle

Lifecycle stage	Key question	Required safeguard	Social doctrine principle
Purpose definition	What human good is served?	Prohibit systems with no legitimate public or human value	Common good
Data and training	Whose lives are represented or erased?	Data provenance, bias testing, privacy, inclusion of vulnerable groups	Justice
Model design	What choices are delegated to the system?	Human agency, explainability proportional to risk, contestability	Dignity
Deployment context	Who is affected and who benefits?	Community consultation and subsidiarity review	Subsidiarity
Monitoring	What harms emerge over time?	Independent audit, incident reporting, redress, and rollback	Solidarity
Distribution of benefits	Who captures value?	Access policies, public-interest licensing, capacity building for poorer nations	Universal destination of goods

The *design* phase must include affected communities. It is not enough for engineers, investors, and regulators to deliberate among themselves. Teachers should help design educational AI; patients and clinicians should shape health AI; workers should shape workplace automation; local governments and communities should shape public-service AI; fragile states and Small Island Developing States should be included in global AI governance because they are often most exposed to digital dependency and least represented in rule-making.

Public procurement can become the main enforcement instrument. Governments, universities, the UN system, development banks, and public agencies should need AI impact assessments before purchase or deployment. These assessments should include dignity, bias, ecological footprint, labor impact, public explainability, cybersecurity, democratic accountability, and grievance mechanisms. The ethical question must become contractual, auditable, and enforceable.

5. From “civilization of love” to technological social justice

The encyclical’s language of the “civilization of love” can appear spiritual, but it has hard policy implications. A civilization of love is not sentimentality; it is a social order that refuses to organize the future around domination. In the AI era, this means rejecting the culture of power in which innovation is valued primarily by market capitalization, military advantage, behavioral capture, or geopolitical leverage.

A civilization of love requires justice in the architecture of technological systems. It asks whether AI contributes to communion or isolation, care or abandonment, truth or manipulation, peace or escalation, inclusion, or exclusion. It calls for the disarming of words in digital spaces and for the rebuilding of dialogue where platforms and political actors' profit from humiliation and antagonism.

Policy translation

- AI for care: prioritize health, education, disaster risk reduction, food systems, climate adaptation, and public services.
- AI for the poor first: build tools for the communities least served by markets, including linguistic minorities, rural populations, fragile states, and SIDS.
- AI for peace: prevent autonomous escalation, algorithmic targeting abuses, and synthetic information warfare.
- AI for truthful communication: support provenance standards, public-interest media, digital literacy, and accountability for manipulation.
- AI for work with dignity: share productivity gains, protect workers, and invest in human capabilities rather than human replacement as an end, in itself.

This is where Christian spiritual life enters the public policy debate without imposing confessional rule. The spiritual claim is that the person is more than utility, more than data, more than market behavior and more than biological optimization. The policy consequence is that AI systems must leave room for conscience, relationship, vulnerability, forgiveness, creativity, contemplation, and responsibility. A society that automates judgment but loses compassion is not advanced. It is impoverished.

6. Convergence with the *One Humanity* framework

The *One Humanity* framework converges with *Magnifica Humanitas* in its core intuition: *humanity* has entered an age of extraordinary power and unprecedented fragility, and global governance must be rebuilt around human dignity and shared survival. In *One Humanity: Geopolitics and the New World Order*, the argument is institutional and geopolitical. The current international system remains fragmented, reactive, and misaligned with the needs of humanity, particularly the poorest nations and Small Island Developing States. The proposed response is an architecture structured around People, Planet, Peace, Prosperity, and Partnership.

Magnifica Humanitas provides the spiritual and ethical foundation; *One Humanity* provides instruments of operationalization. The encyclical asks humanity to choose Jerusalem over Babel. *One Humanity* asks the international system to move from world order to human order. Both reject a future governed by power alone. Both insist that the weakest cannot remain peripheral. Both see technology as a field of moral decision, not an autonomous destiny.

Comparative alignment

Magnifica Humanitas	One Humanity	Combined policy meaning
Human dignity	People-centered geopolitics	AI must serve persons, not systems alone
Common good	One Planet / shared survival	AI must strengthen public goods and planetary resilience
Universal destiny of goods	Equity for poorest nations and SIDS	AI capacity, data, and compute must not deepen global inequality
Subsidiarity	Places and local priorities	AI governance must empower communities and national priorities
Solidarity	Prevention, resilience, and financing	AI transition must protect those most exposed to risk
Civilization of love	Narrative diplomacy and human order	Technology must be framed by responsibility, dialogue, and peace

A decisive *bridge* is the **One Humanity Index**. If adapted to AI, it could measure whether technological systems improve human dignity and resilience across five pillars. Such measurement would convert moral principles into governance indicators, making dignity visible in policy dashboards and investment decisions.

7. Divergences and necessary clarifications

The convergence is strong, but the two frameworks are not identical. *Magnifica Humanitas* is explicitly Christological. It grounds the grandeur of humanity in the Incarnation, the image of God, grace, communion, prayer, and Christian humanism. *One Humanity* is broader and more civic-universal. It speaks across religions, states, cultures, and institutions, building a narrative of planetary solidarity that can include believers and non-believers.

This difference is not a weakness. It is a complementarity. *Magnifica Humanitas* gives depth: an anthropology of the person that resists reduction to machine logic. One Humanity gives reach: a diplomatic and institutional framework capable of translating moral claims into global cooperation, indicators, finance, summits, and reform proposals. The first guards the soul of the project; the second builds its architecture.

A second difference concerns *implementation*. The encyclical is strongest in moral discernment and doctrinal principles. One Humanity is strongest in operational instruments: the One Humanity Index, One Humanity Fund, UN reinvention, One Humanity Diplomacy, and an AI Media and Data Centre. The task is to connect them: not to secularize the encyclical into generic values, but to translate its principles into institutional design while preserving its insistence on the mystery and dignity of the person.

A third clarification concerns “*Humanity*.” In the encyclical, Humanity is the human family called to communion in God and revealed in Christ. In One Humanity, Humanity is the indivisible global subject of shared risk and shared destiny: eight billion people, more than 200 countries and territories, one planet. These meanings are different in source but mutually reinforcing in practice. Both oppose fragmentation, exclusion, and the reduction of persons to instruments of power.

Strategic implication

The policy opportunity is to develop a shared public language: Christian in depth, universal in address, operational in design. This language can speak simultaneously to churches, governments, UN agencies, universities, technology companies, philanthropies, youth movements, and civil society. It can present AI governance not as a defensive regulatory burden, but as a positive construction site for a more human future.

8. Proposed One Humanity AI and Social Justice Compact

This paper proposes a One Humanity AI and Social Justice Compact as the practical bridge between Magnifica Humanitas and the One Humanity framework. The Compact should be short, implementable, and measurable. It should apply to governments, international organizations, universities, technology firms, public agencies, development banks, and philanthropic investors.

Ten commitments

- 1. Human dignity by design: no high-impact AI system should be deployed without a dignity and agency assessment.
- 2. Common-good purpose test: AI systems affecting public life must demonstrate a legitimate social value beyond profit or control.
- 3. Preferential option for the vulnerable: systems must identify and mitigate harms to the poor, migrants, children, workers, minorities, fragile states, and SIDS.
- 4. Human oversight and accountability: critical decisions affecting rights and life chances must remain contestable and accountable to human institutions.
- 5. Truth and communication integrity: synthetic content, automated persuasion and public information systems must be transparent and traceable.
- 6. Work with dignity: automation strategies must include worker consultation, reskilling, social protection, and fair sharing of productivity gains.
- 7. Digital subsidiarity: communities and local authorities should shape AI systems deployed in their contexts.
- 8. Solidarity in access: AI capacity building, language inclusion, open knowledge, and affordable tools should support the Global South.
- 9. Peace and non-weaponization: AI must be governed by human responsibility, arms-control norms and prohibitions on dehumanizing uses.
- 10. Measurement and redress: all high-risk systems should include independent audit, public reporting, accessible complaint channels, and remedies.

Institutional vehicle

The Compact could be hosted through a One Humanity AI Media and Data Centre in partnership with universities, faith-based institutions, UN and regional organizations, civil society, and responsible technology actors. Its first product should be a One Humanity AI Impact Scorecard, aligned with People, Planet, Peace, Prosperity, and Partnership, and explicitly mapped to the social doctrine principles of dignity, common good, universal destination of goods, subsidiarity, solidarity, and justice.

9. Conclusion: *remaining human in the age of intelligent machines*

Magnifica Humanitas and *One Humanity* meet at the same civilizational threshold. The AI era is a test of whether humanity can govern its own power. The encyclical warns that the future can become another Babel: technically impressive, spiritually empty, socially unjust, and politically dangerous. The *One Humanity* framework warns that global governance is already misaligned with the scale and interdependence of the risks humanity faces. Together, they offer a richer answer: rebuild the walls of the common home with dignity at the center, solidarity as method, subsidiarity as governance, justice as measure, and love as the deepest architecture of peace.

The principal lesson is that AI governance cannot be left to markets alone, states alone or engineers alone. It requires a new alliance of conscience and competence: theologians and technologists, diplomats and designers, educators and entrepreneurs, workers and regulators, youth and elders, the Global South and the Global North. Every actor has a section of the wall to rebuild.

For the *One Humanity agenda*, the encyclical provides a powerful moral reinforcement. It affirms that a future human order must be more than institutional reform; it must be a conversion of priorities. Technology must serve life. Power must become responsibility. Innovation must become care. Global governance must protect those most exposed to vulnerability, not only those most capable of influence. The civilization of love and the *One Humanity* framework are therefore not competing visions. They are two languages for the same historical duty: to ensure that the future *is designed* around the dignity of every human being.

Final policy message

The AI era will not be judged by the sophistication of its machines, but by the humanity of its choices. If AI - artificial intelligence *is designed* through dignity, governed for the common good, shared through solidarity, localized through subsidiarity, and measured by justice, it can become part of a new “civilization of love” and a practical “One Humanity” architecture. If not, it will become a new Babel: powerful, efficient, but ultimately *inhuman*.

Selected bibliographic sources:

- (1) Encyclical Letter *Magnifica Humanitas* (2026), Holy See Press Office, on safeguarding the human person in the time of artificial intelligence, May 2026.
- (2) Alexandre Caldas (2026), *One Humanity. Geopolitics and the New World Order*, CIJVS, May 2026.
- (3) Alexandre Caldas (2025), *One Humanity. Time for the Other Side of the World*, Clube de Autores, December 2025.