

Fraternity in the Age of AI

Our Global Appeal for Peaceful Human Coexistence and Shared Responsibility

To His Holiness Pope Leo XIV to all global leaders to all people of goodwill

Moved by a profound desire for a future in which human beings guide society and its choices, we — an independent working group composed of experts, technology leaders, intellectuals, and scholars from different nations, cultures, and faiths — issue this appeal so that artificial intelligence may be developed responsibly, by and for people.

The choices we make today in the field of artificial intelligence will profoundly shape the world we leave to future generations. AI is already causing significant harm: it exacerbates inequalities, concentrates power in the hands of a few actors, and produces damaging environmental effects. Vast and growing sums are being invested in the development of agentic technologies with the potential to surpass human cognitive abilities — what many AI researchers call “superintelligence.” These challenges demand moral leadership and concrete, timely action.

Artificial intelligence offers meaningful opportunities to accelerate scientific research, foster deeper understanding among human beings, revolutionize healthcare, strengthen governance models, and more broadly contribute to building sustainable and inclusive prosperity. However, AI also carries significant risks, as highlighted in the *International Scientific Report on AI Safety*: job displacement, erosion of individual freedoms, power struggles, disinformation and manipulation, mass surveillance, environmental impacts, and broader threats to human well-being.

To fully seize legitimate and potential opportunities while mitigating costs and risks, it is essential to lay the foundations for human prosperity, accompanied by clear and firm boundaries rooted in respect for dignity, community, human and environmental rights, and grounded in the principle of responsibility.

In this spirit of fraternity, hope, and caution, we call upon your leadership to uphold the following principles and non-negotiable limits, with the aim of fostering dialogue and reflection on how artificial intelligence can best serve all humanity.

Principles and Non-Negotiable Limits

- **Human life and dignity**

AI must never be developed or used in ways that threaten, diminish, or deny human life, dignity, or fundamental rights. Human intelligence — our capacity for judgment, moral reasoning, and orientation toward truth and beauty — must never be devalued by any form of artificial processing, however sophisticated.

- **AI must remain a tool, not an authority**

AI must remain under human control. Building uncontrollable systems or excessively delegating decisions is morally unacceptable and must be prohibited by law. Therefore, the development of super intelligent AI technologies must not be authorized until there is broad scientific consensus that they can be built safely and controllably, and until clear public consensus is reached.

- **Responsibility**

Only human beings possess moral and legal status. AI systems are — and must remain — objects of law, never subjects. Responsibilities and obligations fall on developers, providers, companies, implementers, users, institutions, and governments. AI must not be granted legal personality or “rights.”

- **Life-and-death decisions**

AI systems must never be allowed to make life-or-death decisions, especially in military applications during armed conflict or peacetime, in law enforcement, border control, healthcare, or judicial decisions.

- **Safe and ethical development**

Developers must design AI with safety, transparency, and ethics at the center — not as secondary considerations. Users must consider context and potential harms and are subject to the same ethical and safety principles as developers. Independent testing and adequate risk assessment must be required before deployment and throughout the entire lifecycle.

- **Stewardship and accountability**

Governments, companies, and all other actors must not turn AI into a tool of domination, illegal aggression, coercion, manipulation, social scoring, or unjustified mass surveillance.

- **Responsible design**

AI should be designed and independently evaluated to avoid unforeseen and catastrophic effects on humans and society — including systems that generate deception, illusion, dependency, or loss of autonomy.

- **No AI monopolies**

The benefits of AI — economic, medical, scientific, social — must not be monopolized.

- **No devaluation of humanity**

AI design and implementation should help human beings flourish in their chosen activities, not render humanity redundant, displaced, devalued, or replaceable.

- **Ecological responsibility**

Our use of AI must not endanger the planet or ecosystems. AI's vast demands for energy, water, and rare earth materials must be managed responsibly and sustainably across the entire supply chain.

- **No irresponsible global competition**

We must avoid an irresponsible race among companies and nations toward ever more powerful AI.

A Call for Moral Leadership

Respecting these principles will not be easy. It requires moral courage, effective accountability, forward-looking leadership across all sectors of society, and a binding international treaty establishing non-negotiable limits and an independent oversight institution with enforcement and sanctioning powers.

We therefore call for moral leadership for the age of AI.

Because the dangers of AI are often indirect, we urge scientists, civil society, rights groups, and other stakeholders to make greater efforts to articulate — and amplify public awareness of — AI's limits and dangers.

We invite scientists, technology leaders, and policymakers to listen to the voices, experiences, and research of those who work with data, of communities, and of people who bear the consequences of AI's material impacts. We urge them to focus their work on protecting and benefiting the most vulnerable. The legitimacy of moral and legal norms in a society depends on how it treats its most vulnerable members.

We also call on scientists, civil society, and independent auditors to develop and propose new goals and metrics for training, optimizing, and evaluating learning algorithms in terms of truthfulness, balance, and human good, throughout the entire lifecycle — not only in terms of performance and engagement.

We encourage lawmakers, technology leaders, and global communities to collaborate in developing comprehensive guidelines for AI governance that serve the common good. This includes the right of human beings to live without AI. The progress of true human fraternity in the age of artificial intelligence requires the establishment of universal ethical and legal standards.

Finally, we appeal to all people of goodwill: let us unite to ensure that AI serves all humanity rather than a narrow minority.

By joining across nations, cultures, and beliefs, and by prioritizing dialogue over competition, we can shape a future that elevates human dignity and promotes a more just and peaceful world.

We invite all stakeholders — citizens, scientists, business leaders, religious leaders, community representatives, and policymakers — to participate in this initiative. Together, we reaffirm the essential principle that machines must serve the interests of humanity.