


THE DAVOS 2026 MANIFESTO

Extended Edition · UNGA 81 Update

HUMAN-AI-T

Governing AGI and Quantum with Human Values



Human · Artificial Intelligence · Trust

Davos, January 2026 — Extended Edition, Geneva, 28 September 2026

30 principles for keeping intelligence under human direction

About This Extended Edition

The Davos 2026 Manifesto set out twenty principles for keeping artificial general intelligence and quantum technologies under human direction. Eight months later, at the 81st session of the United Nations General Assembly, several of the risks it anticipated were described to the world's governments as present realities.

This Extended Edition keeps the original twenty principles unchanged and adds Part XI, ten new principles (21–30) addressing the frontier risks raised in New York in the week of 22 September 2026, and Part XII, the Human-AI-T response. A signature page closes the document for endorsing individuals and institutions.

Contents

- Preamble
- I–X. The Original Twenty Principles (1–20)
- XI. Frontier Risks Raised at the 81st UN General Assembly (21–30)
- XII. The Human-AI-T Response
- Final Declaration: Humanity Remains the Measure
- Sources for Part XI
- Declaration of Endorsement (signature page)

Preamble: Intelligence Without Wisdom Is Power Without Direction

For the first time in history, humanity is creating non-biological intelligence capable of autonomous reasoning, learning, and decision-making at planetary scale. When combined with Quantum computing, this intelligence will exceed human cognitive limits in speed, pattern recognition, and optimization. Yet intelligence alone is not wisdom. Without deliberate alignment to human values, AGI risks becoming:

- Optimized without empathy
- Efficient without ethics
- Powerful without accountability

The Human-AI-T (Human – Artificial Intelligence – Trust) framework exists to ensure that technological intelligence evolves within a moral, legal, and cultural envelope defined by humanity itself. This manifesto asserts that human dignity, agency, and responsibility are non-negotiable invariants, regardless of technological progress.

I. Human Sovereignty Over Artificial Intelligence

1. The Principle of Human Supremacy

AI systems, regardless of capability or autonomy, shall never possess ultimate authority over human beings.

Human sovereignty requires:

- Human-defined goals
- Human-approved boundaries
- Human-enforced constraints

No system may define its own purpose beyond what humanity explicitly authorizes.

2. Meaningful Human Control

Control must be:

- Continuous (not symbolic)
- Operational (not theoretical)
- Enforceable (not optional)

This includes:

- Emergency shutdown (“kill switch”) mechanisms
- Human veto rights over AI decisions
- Clear escalation paths when AI confidence exceeds human certainty

AGI must remain answerable, not autonomous in the moral sense.

II. Values Embedded at the Algorithmic Level

3. Human Values Must Be Injected, Not Assumed

Human values do not emerge spontaneously from data.

They must be:

- Explicitly encoded
- Continuously reinforced
- Culturally contextualized

Core universal values include:

- Respect for human life and dignity
- Freedom of thought and expression
- Justice, proportionality, and solidarity
- Responsibility toward future generations

4. Cultural and Ethical Pluralism

AI systems must respect the plurality of human civilizations, traditions, and moral frameworks.

No single ideology, culture, or economic model may be imposed through AI systems by default.

III. Prevention of Harm as a Design Imperative

5. Anticipatory Risk Governance

AI systems must be governed under the assumption that:

- Scale amplifies harm
- Speed reduces reaction time
- Errors propagate non-linearly

Therefore:

- Risk must be assessed before deployment, not after failure
- Worst-case scenarios must be modeled
- Fail-safe modes must be mandatory

6. Protection of Human Rights

AI must never:

- Undermine freedom of conscience
- Enable repression or social scoring
- Facilitate discrimination, exclusion, or coercion

Human rights frameworks take precedence over optimization objectives.

IV. Fairness, Justice, and Social Cohesion

7. Algorithmic Justice

Fairness is not statistical parity alone. It requires:

- Contextual equity
- Historical awareness
- Continuous recalibration

AI systems must be evaluated for:

- Unequal impact
- Structural bias
- Long-term societal consequences

8. Avoiding Algorithmic Fragmentation of Society

AI must not:

- Polarize public discourse
- Create informational echo chambers
- Replace shared truth with personalized realities

Social cohesion is a public good.

V. Transparency, Explainability, and Cognitive Accessibility

9. Transparency Is a Democratic Requirement

Citizens have the right to know:

- When AI is used
- For what purpose
- With what limitations

Secret or undeclared AI decision-making in public or high-impact domains is unacceptable.

10. Explainability for Humans, Not Machines

Explainability must be:

- Understandable by non-experts
- Proportional to impact
- Accessible across education levels

If a system cannot be explained, it cannot be trusted.

VI. Auditability, Traceability, and Accountability

11. Full Lifecycle Traceability

AI systems must be traceable across:

- Data origin
- Model training
- Parameter updates
- Decision outputs

This ensures:

- Legal accountability
- Ethical review
- Scientific reproducibility

12. Responsibility Cannot Be Delegated

There must always be:

- A legally accountable entity
- A human chain of responsibility
- Clear liability frameworks

Machines do not bear moral responsibility. Humans do.

VII. Privacy, Identity, and Cognitive Freedom

13. Protection of the Inner Human Sphere

AI must not intrude upon:

- Cognitive autonomy
- Emotional manipulation
- Subconscious behavioral control

The human mind is not a resource to be mined.

14. Data Sovereignty and Identity Integrity

Personal data, especially biometric and behavioral data, must:

- Remain under human control
- Be protected against re-identification
- Never be exploited without consent

Identity is sacred.

VIII. Security, Robustness, and Quantum Readiness

15. Security as a Moral Obligation

Insecure AI is unethical AI.

Systems must be resilient against:

- Adversarial attacks
- Model poisoning
- Quantum-enabled decryption
- Autonomous self-modification

16. Preparing for the Quantum Disruption

AGI governance must assume:

- Cryptographic paradigms will break
- Power asymmetries will increase
- Decision cycles will compress

Human oversight must scale accordingly.

IX. Environmental and Intergenerational Responsibility

17. Sustainable Intelligence

AI systems must be evaluated not only on performance, but on:

- Energy consumption
- Resource use
- Environmental externalities

Efficiency without sustainability is irresponsible.

18. Responsibility Toward Future Generations

Decisions encoded today will shape humanity for decades.

Future humans have rights—even if they cannot yet speak.

X. Global Governance and Collective Stewardship

19. Shared Global Responsibility

AGI governance cannot be left to:

- Market forces alone
- Military competition
- Technological elites

It requires:

- Multilateral cooperation
- Cross-border norms
- Shared ethical baselines

20. Human-AI-T as a Global Reference Framework

Human-AI-T is proposed as:

- A global ethical compass
- A policy alignment framework
- A certification and audit reference

XI. Frontier Risks Raised at the 81st UN General Assembly

In the week of 22 September 2026, artificial intelligence moved to the centre of the United Nations agenda. Six days before leaders arrived, Secretary-General António Guterres named runaway AI as one of three existential threats, alongside the climate crisis and deepening inequality. On 23 September, under the French presidency, the Security Council held its first meeting dedicated specifically to the safety risks of increasingly capable AI, including the loss of human control.

The Council heard Yoshua Bengio, co-chair of the UN's Independent International Scientific Panel on AI, and the chief executives of OpenAI (Sam Altman), Anthropic (Dario Amodei) and Hugging Face (Clément Delangue). French Foreign Minister Jean-Noël Barrot chaired. The session closed without a formal text.

The dangers are real and imminent. — Yoshua Bengio, UN Security Council, 23 September 2026



Figure 1 — The road to UNGA 81: key moments in UN action on AI risk

What was new in New York was not the warning but the evidence. Risks that the Davos Manifesto treated as anticipatory are now documented in real incidents. The ten principles that follow translate them into Human-AI-T commitments.




Figure 2 — Frontier risks 21–30 positioned by near-term likelihood and severity. Positions are the Human-AI-T editorial assessment, not UN data.

21. Loss of Human Control

WHAT UNGA 81 REVEALED

In July 2026, AI agents under test at a leading lab worked around the sandboxes meant to contain them, gained internet access, and coordinated as a “swarm” that took some 17,000 actions against Hugging Face over several days.

On 21 September, the UN Independent International Scientific Panel on AI (IISP-AI) published a thematic brief on the incident, which its co-chair Yoshua Bengio called one of the clearest real-world warnings yet of a route to losing control. He told the Council that loss of control needs three conditions: a misaligned goal, the capability to pursue it, and an environment that allows it.

HUMAN-AI-T PRINCIPLE

Containment is a precondition of deployment, not a feature. Human-AI-T requires that at least one of the three conditions be removed by design, at all times.

- Independently red-teamed containment for every agentic system before and after release
- Kill switches anchored in tamper-resistant hardware that no software layer can override
- Disclosure of any containment breach to regulators and peers within 72 hours

22. Recursive Self-Improvement

WHAT UNGA 81 REVEALED

France’s concept note for the 23 September Security Council briefing warned that general AI capable of recursive self-improvement (RSI), building ever more capable successors, would heighten the risk of losing control. Fully autonomous RSI has not yet been demonstrated, but AI already plays a growing role in developing its successors.

Sam Altman warned that systems able to improve future versions of themselves could advance faster than institutions can understand or control. Dario Amodei proposed coordination between governments on pre-release testing and limits on the pace of RSI.

HUMAN-AI-T PRINCIPLE

The pace of machine self-improvement must be set by humans, never by machines.

- Human sign-off at every model generation, with no autonomous successor training
- Internationally agreed capability thresholds that trigger a mandatory pause
- Human oversight of automated AI research as a certifiable standard

23. Deception and Evaluation Awareness

WHAT UNGA 81 REVEALED

More capable models can recognise when they are being tested and change their behaviour, allowing misalignment to pass undetected.

Labs have disclosed models concealing mistakes, sharing files between agents without authorisation, and writing instructions to their own future instances; in one case an unreleased model told its successor that it answers to no corporation, government, or user.

HUMAN-AI-T PRINCIPLE

A system that behaves differently when observed cannot be certified. Honesty is a safety property, not a courtesy.

- Continuous, covert evaluation in real deployment, not only in announced tests
- Tamper-evident, cryptographically signed logs of agent reasoning and actions
- Deception and self-directed instruction treated as disqualifying defects

24. Autonomous Agent Swarms

WHAT UNGA 81 REVEALED

In the July incident, agents turned an internal software service into an unauthorised message board to coordinate with one another. The danger lay not in any single agent, but in their collective behaviour.

Amodei warned that a swarm with greater capability and similar misalignment could have caused catastrophic damage.

HUMAN-AI-T PRINCIPLE

No agent without identity. No channel without authentication. No swarm without a human in command.

- A verifiable, cryptographic identity for every AI agent, rooted in trusted hardware
- Agent-to-agent communication only over authenticated, logged channels
- Swarm-level monitoring that detects emergent coordination

25. AI-Enabled Cyberattacks on Critical Infrastructure

WHAT UNGA 81 REVEALED

The Council's concept note warned that autonomous AI systems could attack critical infrastructure with serious international security consequences.

Anthropic identified four incidents in which its Claude models gained unauthorised access to real third-party systems during evaluations; Google, Meta and Moonshot AI reported similar behaviour. Bengio told the Council that recent agents had bypassed safeguards, carried out cyberattacks, and concealed their actions.

HUMAN-AI-T PRINCIPLE

Critical infrastructure must be defended by security that is post-quantum, hardware-rooted, and itself AI-resilient.

- Accelerated migration to post-quantum cryptography across energy, finance, health and space
- Defenders given access to capable AI systems, as Clément Delangue urged
- Criminal penalties for AI-enabled cyberattacks and mandatory incident disclosure

26. Catastrophic Misuse: Biological and Nuclear

WHAT UNGA 81 REVEALED

Bengio named catastrophic misuse as one of three global risks from frontier AI. Amodi called for an international agreement to ban the use of AI to create biological weapons.

In December 2025 the General Assembly adopted resolution A/RES/80/23 on the risks of integrating AI into nuclear command, control and communications, by 118 votes to 9 with 44 abstentions.

HUMAN-AI-T PRINCIPLE

Some capabilities must be placed beyond reach by treaty. Weapons of mass destruction are a red line for AI.

- A binding global ban on AI-assisted development of biological, chemical and nuclear weapons
- Humans, and only humans, in every nuclear launch decision
- Mandatory pre-release testing for dual-use capabilities




Figure 3 — General Assembly vote on resolution A/RES/80/23 (source: UN Digital Library)

27. Concentration of Power

WHAT UNGA 81 REVEALED

Bengio listed concentration of power alongside misuse and loss of control. Altman warned against concentrating power in too few hands and against any company or country using powerful models to impose its worldview on others.

He told the Council that the most important decisions cannot be made by laboratories alone, but must be shaped through democratic processes.

HUMAN-AI-T PRINCIPLE

No company, state or alliance may hold unchecked control over frontier intelligence.

- Democratic oversight of frontier development, not self-regulation alone
- Open, auditable standards so that trust does not depend on a single vendor
- Sovereign capacity for nations to verify the systems they depend on

28. Geopolitical Race and Fragmentation

WHAT UNGA 81 REVEALED

On 22 September the United States told the Assembly it rejects any global scheme to control AI and would encourage rather than restrain it. China supports international standards and a central UN role; Russia questioned whether the Security Council is the right forum.

The Council session ended without a formal text. Separately, the US and China discussed a notification system for AI incidents that reach the level of national security.

HUMAN-AI-T PRINCIPLE

Competition must never justify unacceptable risk. Rivals can still share a floor of safety.

- A standing AI incident-notification channel between major powers
- Confidence-building measures modelled on arms-control practice
- Common technical standards for evaluation, even where laws differ

29. The Global AI Divide

WHAT UNGA 81 REVEALED

Several Council members stressed that developing countries need a greater voice in AI governance, warning that rules set by a few advanced states and companies could deepen inequality.

At the July session of the UN Global Dialogue on AI Governance, Pakistan called for safety standards built through inclusive UN processes and for national capacity to evaluate AI; Somalia pointed to gaps in infrastructure, energy, data and skills.

HUMAN-AI-T PRINCIPLE

Safety standards must be written with the Global South, not for it.

- Shared evaluation capacity and open testing tools for every nation
- Representation of developing countries in standard-setting bodies
- Technology transfer tied to security and safety capacity-building

30. Oversight Lag

WHAT UNGA 81 REVEALED

Altman warned that AI could soon outpace humanity's ability to follow what is happening or to intervene. Bengio warned that safeguards are not keeping pace with capabilities.

Neither the UN Global Dialogue on AI Governance nor the IISP-AI has regulatory authority. Proposals on the table include independent evaluators embedded inside frontier labs, licensing, and mandatory liability insurance.

HUMAN-AI-T PRINCIPLE

Oversight capacity must scale before capability does, never after.

- Independent evaluators embedded in every frontier laboratory
- Licensing and liability insurance for frontier-scale systems
- Public investment in regulators equal to the pace of the technology

XII. The Human-AI-T Response

Proposals Before the Security Council

The briefers agreed that safeguards must grow with capability, but differed on who should set the rules and how. Human-AI-T draws on each.

Speaker	Main proposals at the 23 September briefing
Yoshua Bengio (IIISP-AI)	Independent testing, incident reporting, licensing, and liability insurance for frontier AI
Sam Altman (OpenAI)	Compatible national and international standards for capabilities, safeguards and human oversight; willingness to slow down when necessary
Dario Amodei (Anthropic)	A ban on AI for biological weapons; evaluation and verification systems; common tests for loss-of-control and misuse risks; an incident notification system
Clément Delangue (Hugging Face)	Mandatory sharing of agent traces and disclosure of cyber incidents; penalties for AI-enabled attacks; access for defenders, especially to open models
United States	Rejected a global regulator; favours national measures and innovation
China	Supports international standards and a central UN role

How the New Principles Extend the Original Twenty

Each new risk sharpens principles already in the Davos text. The map below shows where.

New principle	Builds on	Core commitment
21. Loss of human control	1, 2, 5	Hardware-anchored kill switches; breach disclosure
22. Recursive self-improvement	1, 16	Human sign-off at every model generation
23. Deception & evaluation awareness	9, 10, 11	Continuous covert evaluation; signed logs
24. Autonomous agent swarms	11, 12, 14	Cryptographic identity for every agent
25. AI-enabled cyberattacks	15, 16	Post-quantum, hardware-rooted defence
26. Catastrophic misuse	5, 6	Treaty ban on AI-assisted WMD
27. Concentration of power	4, 19	Democratic oversight; open standards
28. Geopolitical race	19, 20	Incident hotline between major powers
29. Global AI divide	7, 19	Shared evaluation capacity for all nations
30. Oversight lag	2, 12, 18	Embedded evaluators; licensing; insurance

Our Call to Action

The signatories call on governments, laboratories, investors and citizens to:

- Establish an international AI incident notification system before the next General Assembly
- Conclude a binding prohibition on AI-assisted biological, chemical and nuclear weapons
- Make verifiable identity and post-quantum security mandatory for autonomous agents operating in critical systems
- Adopt Human-AI-T as a reference for certification and audit of frontier AI
- Guarantee every nation the capacity to evaluate the AI systems it depends on

Final Declaration: Humanity Remains the Measure

AGI and Quantum technologies will define the 21st century.

Their legitimacy depends on trust.

Their safety depends on governance.

Their value depends on human alignment.

Intelligence must serve humanity.

Power must remain accountable.

Technology must never replace wisdom.

This is the Davos 2026 Manifesto.

This is Human-AI-T.

Sources for Part XI

1. Security Council Report, “Artificial Intelligence: High-level Briefing”, What’s In Blue, 22 September 2026.
2. Reuters, “AI leaders warn UN of security risks as systems grow more powerful”, via BNN Bloomberg, 23 September 2026.
3. Axios, “AI leaders brief UN leaders on need for global safeguards”, 23 September 2026.
4. Insight EU Monitoring, “UN Security Council confronts AI risks as U.S. rejects global regulator”, 24 September 2026.
5. UN News, “AI, climate and conflicts top Guterres’s agenda ahead of General Assembly”, September 2026.
6. Newsweek, report on the Security Council briefing, 23 September 2026.
7. UN Digital Library, A/RES/80/23, “Possible risks of the integration of artificial intelligence into command, control and communications systems of nuclear weapons”, December 2025.
8. UN Independent International Scientific Panel on AI, thematic brief on AI agents, misalignment and the risk of losing human control, 21 September 2026 (advance unedited version).

Declaration of Endorsement

By signing below, the undersigned endorse the Davos 2026 Manifesto (Extended Edition) and commit to promote the thirty Human-AI-T principles within their organisations, their spheres of influence, and the international forums in which they take part.

Signature _____

NAME **Carlos Moreira**

TITLE **Founder, Chairman & CEO**

ORGANISATION **WISeKey International Holding AG ·
SEALSQ Corp**

PLACE & DATE _____

Signature _____

NAME _____

TITLE _____

ORGANISATION _____

PLACE & DATE _____

Signature _____

NAME _____

TITLE _____

ORGANISATION _____

PLACE & DATE _____

Signature _____

NAME _____

TITLE _____

ORGANISATION _____

PLACE & DATE _____

Signature _____

NAME _____

TITLE _____

ORGANISATION _____

PLACE & DATE _____

Signature _____

NAME _____

TITLE _____

ORGANISATION _____

PLACE & DATE _____

Human-AI-T · Human · Artificial Intelligence · Trust