

PUBLIC MANIFESTO

Milan, global capital of artificial intelligence governance

A commitment by the Milanese and Lombard system to establish in Milan a new United Nations pillar dedicated to the governance of artificial intelligence.

A manifesto originating from the work of the American Chamber of Commerce in Italy and open to the endorsement of the Milanese and Lombard system.

On 26 December 1883, at the premiere of Ponchielli's *La Gioconda*, the Teatro alla Scala lit up with electric light thanks to the Santa Radegonda power station, the first in continental Europe, conceived by Giuseppe Colombo, an engineer of the Politecnico, with Thomas Edison's technology. From Milan, across a transatlantic bridge, electricity entered continental Europe. 143 years later, artificial intelligence is once again arriving from the United States and knocking at the same door.

The United Nations is building its own architecture for AI governance. General Assembly Resolution A/RES/79/325 of August 2025 established the Independent Scientific Panel and the Global Dialogue. The first session of the Global Dialogue was held in Geneva on 6 and 7 July; the second is scheduled for New York in May 2027. The United Nations AI Governance for Humanity Lab, inaugurated in Valencia in May 2026, works on practical cooperation among stakeholders. UNESCO, from Paris, oversees the normative dimension with its 2021 Recommendation on the Ethics of AI. There is still no structure dedicated to the horizontal coordination of AI governance that would give the system common evaluation methods. That institutional space will be defined over the next twenty-four months.

The first institutional checkpoint is already behind us. On 4 and 5 September the Pontifical Gregorian University hosted the roundtable AI and Human Development, convened by the Deputy Secretary-General of the United Nations, whose outcomes will contribute to the implementation of the Pact for the Future and support the Global Dialogue. On 21 September the Independent Scientific Panel published its first thematic brief, devoted to the risk of losing control of autonomous agents. During the high-level week of the 81st General Assembly, from 22 to 28 September, AI governance was at the centre of the debate. Foreign Minister Antonio Tajani called for common rules but not bureaucratic cages, stating that defining them is precisely the task of the United Nations, before events or the strongest impose them.

On 28 July 2026, 1,386 employees of the leading frontier laboratories signed the letter “Pacing the Frontier”, asking the US government to support an international effort on the technical and governance tools needed to slow down, if and when necessary, the race of automated AI. The signatories include senior leaders of OpenAI and Anthropic, starting with Dario Amodei. It is the very researchers who build AI who are calling for international governance of the system. The urgency is no longer an opinion; it is a fact.

In September the request has taken on the voice of those who lead those laboratories. In the essay “We Must Pace the Frontier”, published in September, Anthropic co-founder and CEO Dario Amodei wrote that “we must slow the pace at which we improve the capabilities of AI models”, citing as causes the acceleration of recursive self-improvement of systems and recent episodes of agentic misalignment, and estimating at six to twelve months the window within which a sufficiently capable swarm of agents could compromise significant parts of the network. The proposal is structured in three steps: third-party evaluators embedded in the laboratories with continuous access, coordination among democracies on common safety standards, and global coordination leading to a shared “speed limit” on recursive self-improvement and to verifiable agreements with the People's Republic of China. “We owe it to humanity to try,”

Amodei concludes. Two of the three steps are multilateral by definition, and today they have no institutional home in which to be negotiated, verified and legitimised. That is exactly the home that Milan offers to build.

On 2 September 2026, in Chapel Hill, the G20 Innovation Ministers adopted by consensus the Carolina Principles for Emerging Technologies, also endorsed by the People's Republic of China. The framework promotes a sectoral, proportionate and adaptable approach, oriented towards human-centred innovation and the safe and trustworthy adoption of AI, while avoiding the creation of new generalist bureaucracies. The heads of the leading frontier laboratories took part at the same table, including Anthropic co-founder Tom Brown. The G20 framework completes an architecture on the side of the industrial and regulatory policy of States; it does not cover the properly multilateral level or the dimension of ethical principles and fundamental rights, which remain within the remit of the United Nations. It is precisely in that complementary space that the Milanese pillar is positioned.

On 29 September 2026, at the White House, the President of the United States Donald Trump and the heads of Google, Meta, Anthropic, OpenAI, Nvidia and xAI signed the White House Accord on Super Intelligence, a “Joint Commitment on Frontier Responsibilities”. The Accord, voluntary and without binding legal force, commits each company that trains and releases frontier models to four layers of control and verification: internal checks on the capabilities and alignment of models, with particular attention to cybersecurity, biosecurity and chemical threats; an internal team to ensure that they function; independent evaluations entrusted to an external auditor or evaluator; and an independent board committee to oversee their results. The text recognises that over time it may be appropriate to translate these steps into laws or regulations. It is the laboratories that build the frontier, together with the government of the country that hosts it, that recognise that independent evaluation is the condition for trust: the first of the three steps indicated by Amodei thus enters the companies' commitments. By its nature, the Accord remains a national and voluntary commitment among companies, and leaves open the question of how to make comparable, across different countries, the methods by which that evaluation is conducted. It is to this question that the functions of the Milanese pillar respond, measuring without regulating, in full respect of the sovereignty of States.

Milan has the credentials to occupy that space. Assolombarda estimates growth of 1.2% in the city's GDP in 2026. In 2024 private equity investment in Lombard companies approached €9 billion, 60% of the national total. 7 data centres in the Milan area have been recognised as being of national strategic interest. Bending Spoons, listed on Nasdaq in July, is today worth more than \$20 billion. Politecnico, Statale, Bicocca, Bocconi, Human Technopole, Mario Negri and CNR make up a research ecosystem of European critical mass. The Vatican, one hour away by air, has offered, with the Note “Antiqua et nova” of January 2025 and with Leo XIV's encyclical “Magnifica humanitas”, published on 25 May 2026, a moral anchor that no other AI capital can claim.

Milan and Lombardy are also an example of a research ecosystem joined to a manufacturing heart. This is a central point, because the phase now opening no longer concerns software alone. Physical AI, artificial intelligence that perceives its environment, decides and acts in the real world through robots, autonomous machines, and intelligent systems and infrastructure, is set to have a profound impact on people, societies and businesses. Milan is a natural place to experiment with this transition to intelligence that acts in the real world.

Alongside this is the city's ongoing experience as an urban laboratory for responsible AI. The Municipality of Milan has promoted a Manifesto for the use of artificial intelligence, a pioneering experience among European cities, consistent with the principles of the AI Act, and has mapped 45 AI use cases already in production in the Administration and in city services. The Smart City Lab and the Human AI+ programme, now in its second edition, translate those principles into experiments and businesses, while the Board for Technological Innovation and Digital Transformation brings together the public and private sectors, universities and the third sector in an open governance method. Milan does not offer an empty seat. It offers an ecosystem in which the principles of responsible AI are already being tested in the complexity of a large European city.

The EMA precedent of 2017 teaches us that the best of technical dossiers, with Milan leading in the first two rounds of voting and defeated only by the draw after the tie in the third, can lose in the last mile if the country system is not united and if bilateral relations with the decisive Member States do not start in time. We will not repeat that mistake.

There is a coincidence that cannot be overlooked. The architecture of multilateralism was born in San Francisco on 26 June 1945 with the signing of the Charter of the United Nations. Eighty years later, San Francisco is the world capital of frontier artificial intelligence, a revolution that is not only industrial but cognitive. The same place that saw the birth of the UN is today witnessing the birth of the technology that, more than any other, challenges the international order of the twentieth century. Both respond to the same question: how to govern a discontinuity that no State can face alone.

This is where the Sputnik Moment of the United Nations begins. At a time of profound rationalisation of the Organization, the governance of artificial intelligence is not one more agency to be added to the dozens of existing agencies, funds and programmes. It is the horizontal vector around which the UN regains its historic centrality, the new-generation pillar, distinct from the twentieth-century agency model, that gives new meaning to the universalism of the United Nations in the twenty-first century. It is within this logic of rationalisation, consistent with the UN80 method and with the light-touch spirit of the Carolina Principles, that the pillar does not duplicate existing agencies or introduce new regulatory burdens, but connects them and offers the horizontal plane that the system lacks today.

The pillar's line of responsibility runs along the United Nations' own perimeter. Operational support to the Global Dialogue on AI Governance and to the Independent Scientific Panel established by Resolution A/RES/79/325. Interoperability among the regulatory frameworks of States, in continuity with the sectoral and adaptable approach recognised by the G20. Capacity building for the Global South, so that the rules that concern everyone are also written by those who do not sit at the table today. Independent evaluation of public-sector AI pilots and of their impact on work and services. Stewardship of the ethical dimension and of fundamental rights, rooted in the Magisterium of the Catholic Church and in the European constitutional tradition, which no other international table explicitly covers.

The twelve functions of the pillar

AmCham Italy's Policy Paper II, "Milan and the future of global governance of artificial intelligence. The functions of the new UN pillar for AI governance", translates this line of responsibility into twelve technical functions. They are not an invention: they make operational the principles of the UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted in 2021 by all Member States, within the framework of Resolution A/RES/79/325, and they are ordered by a single criterion: technology must remain at the service of the person. Each has a declared limit and a result indicator, and is exercised together with the organisation competent for the subject matter, which retains ownership of it.

01

Pace of the frontier and early warning

Measuring with a public method how fast the capabilities of the most advanced systems are growing, and alerting States in good time to those with systemic implications.

02

Impact assessment

A common method, updated over time, to assess the effects of a system on people.

03

Human responsibility

Whoever is accountable for a decision always remains identifiable along the technological chain.

04

Transparency and provenance

Distinguishing the authentic from the synthetic and obtaining explanations proportionate to the decision.

05

Data and documentation

Documenting what data a system was trained on and on what basis.

06

Equity among countries

A comparable index of gaps in compute, data and skills, in order to reduce them.

07

Work and transition

Comparable data on employment effects and on the transition policies that work.

08

Health and minors

Comparable evidence and precautionary criteria for the most vulnerable people.

09

Training of regulators

Skills for those who regulate and those who use, a precondition for all the other functions.

10

Languages and pluralism

Measuring which languages and cultures systems represent, and supporting those that are excluded.

11

Environment and compute

Making comparable the energy, water and land footprint of data centres.

12

Safety and stability

A common language to recognise and report incidents and dangerous capabilities.

The functions generate four products: methods, measurements, warnings, skills · No normative, authorising or sanctioning power

None of the four products is a rule and none entails compliance obligations. The pillar does not authorise or register models, does not supervise or sanction companies, does not produce horizontal rules, and does not take competences away from ILO, WHO, UNESCO, UNDP, UNCTAD, ITU, WIPO and OHCHR, nor the scientific role from the Independent Panel. Implementation proceeds in three phases: within twelve months, the functions that do not require new mandates; from twelve to thirty-six months, consolidation; beyond thirty-six months, the functions that require scale. It is the space in which both those who call for less regulation and those who call for more verification can come together: measuring does not mean prohibiting; it means making every subsequent choice possible.

The positioning of the Milanese bid is today distinctive. Since last March, the Republic of Korea has been pursuing the physical aggregation into a single campus of the AI functions of six existing UN agencies, a twentieth-century agency model. On 16 July China launched in Shanghai the World AI Cooperation Organization, an alternative intergovernmental organisation signed by twenty-nine States and placed outside the UN perimeter. Kazakhstan has offered Astana as the venue for the first meeting of that body. On 21 September, on the sidelines of the General Assembly, Finland and Norway launched “A Call for Control of Frontier AI Models”, which asks UN Member States to explore an international institution able to set standards, enable verification and convene States when capability thresholds are crossed. The call, welcomed by Secretary-General Guterres, has been signed by 29 countries and the European Commission, including Germany, France, Spain, the Netherlands, Greece, Portugal, Canada and Australia; Italy, the United States, the United Kingdom and China are not among the signatories. Milan is bidding to host a new-generation pillar within the perimeter of the United Nations, with a spontaneous bid that comes from civil society and business even before it comes from a government. It is a positioning still without precedent.

The governance of artificial intelligence cannot be reduced to a table among the technologically dominant economies alone, with Europe and the Global South as spectators of the rules that concern them. Not in opposition to the United States and China, essential and complementary interlocutors in any credible multilateral framework, but because the horizontal and universal nature of the technology requires an inclusive table. Milan is bidding to be the European seat of that table, where the method of law, fundamental rights and institutional pluralism, already codified by the AI Act, is added as a third constitutive axis of the global governance of artificial intelligence. It is the space in which Europe can once again compete, build and grow on the ground that, more than any other, will define the continent's competitiveness in the coming decades, restoring to the Union an institutional centrality that the twentieth-century agency model had not granted it.

The most innovative feature of the proposal concerns its financial structure. The pillar would rest on a hybrid three-tier model, comprising assessed contributions from Member States, permanent voluntary contributions from the main operators in the sector following the successful precedent of GAVI and CEPI in the health field, and thematic partnerships with philanthropic foundations and international financial institutions. The model reduces pressure on the Organization's regular budget and offers a useful institutional precedent for the financial reform of the UN system.

For these reasons, the undersigned commit to

- **Publicly and institutionally support** — Milan's bid to host a new United Nations pillar dedicated to the governance of artificial intelligence, set out in the policy paper “Milan, global capital of ethical artificial intelligence” presented by AmCham Italy on 27 July 2026 and developed in the twelve functions of Policy Paper II of September 2026.
- **Establish a permanent Promoting Committee** — with an operational seat in Milan, bringing together the business, commercial, academic, scientific, professional and financial representatives of the Milanese and Lombard area, and maintaining a continuous dialogue with the Presidency of the Council of Ministers, the Farnesina, MIMIT, the Lombardy Region and the Municipality of Milan.
- **Contribute their own expertise** — to the technical dossier of the bid, each according to its institutional mission, making available research, data, sectoral experience and the capacity for dialogue with international stakeholders, as well as making available to the dossier the initiatives through which the Milanese and Lombard productive system is already testing and industrialising artificial intelligence applied to the physical world.
- **Promote diplomatic action** — by urging the Italian Government and all the political forces of the constitutional spectrum to make the bid a foreign policy priority for the country with the United Nations Secretariat, the Permanent Mission of Italy in New York and the decisive Member States, from the group of co-facilitators Spain and Costa Rica to the leaders of GRULAC, Africa and the G77, ahead of the Global Dialogue on AI Governance in New York on 3 and 4 May 2027.
- **Ensure consistency with the G20 architecture** — in dialogue with the major frontier laboratories and the competent international organisations, so that the Milanese pillar presents itself as a light structure, complementary to the Carolina Principles of 2 September 2026 and to the voluntary commitments of the White House Accord on Super Intelligence of 29 September 2026, avoiding any duplication and qualifying as the UN seat of the ethical, human-centred and independent technical capacity dimension that the G20 framework does not cover.
- **Recognise the urgency of the time window** — The useful window for positioning Milan is the next twenty-four months. Whoever comes to that table with a ready seat, a credible technical dossier, a sustainable financial structure and a cohesive political coalition will shape the institutional geography of AI governance for the next generation. Whoever arrives later negotiates conditions written by others.

The meaning of this signature

We are not signing a petition. We are signing an operational commitment, with the same seriousness with which, in 2017, the Milanese system, with Diana Bracco on its steering committee, supported Milan's EMA bid, and with the added awareness that what is at stake today is the governance of the technology that will redefine the economy, work, information and democracy over the next twenty years.

Every signatory is given equal weight. AmCham Italy, Assolombarda and Confcommercio Milano Lodi Monza Brianza promote the Manifesto, and the Promoting Committee is born civic, cross-cutting and free of partisan labels. No logo prevails over the others, no institution speaks on behalf of the others. The dossier belongs to Milan, not to a single house.

Milan is not asking for a privilege. It offers the country system a concrete opportunity to play a leading role in building the international architecture that will govern artificial intelligence. It also offers a place in which AI governance can maintain a direct relationship with its application. It also belongs to an Italian tradition that has already given the world one of the first shared commitments on the ethics of artificial intelligence: the Rome Call for AI Ethics, promoted by the Pontifical Academy for Life and signed in Rome on 28 February 2020 by Microsoft, IBM, FAO and the Italian Government. Effective governance will help create the conditions for artificial intelligence to be used safely, responsibly and positively to generate innovation, well-being and growth. The proposal is on the table. Let it be discussed at the levels it deserves, with the urgency and seriousness that the moment demands.

Alea iacta est.

SIGNATORIES

Endorsement of the Manifesto

PROMOTERS

American Chamber of Commerce in Italy

Simone Crolla, Managing Director · Stefano Lucchini, President

Assolombarda

Alvise Biffi, President

Confcommercio Milano Lodi Monza Brianza

Carlo Sangalli, President

SIGNATORIES

Claudio Bassoli

President & CEO, HPE

Elena Beccalli

Rector, Università Cattolica del Sacro Cuore

Francesco Billari

Rector, Bocconi University

Giovanni Bozzetti

President, Fondazione Fiera Milano

Marina Brambilla

Rector, University of Milan

Vincenzo Di Nicola — in a personal capacity

Senior Partner, Artificial Intelligence Fund, CDP Venture Capital

Fiorenzo Marco Galli

Director General, Museo Nazionale Scienza e Tecnologia Leonardo da Vinci

Valentina Garavaglia

Rector, IULM University

Teodoro Lio

CEO, Accenture Italia

Daniela Mainini

President, Centro Studi Grande Milano

Emma Marcegaglia

Chair and Chief Executive Officer, Marcegaglia Holding

Giorgio Medda

CEO, Azimut

Letizia Moratti

entrepreneur and philanthropist, former Mayor of Milan

Msgr. Vincenzo Paglia

President Emeritus, Pontifical Academy for Life

Antonio Palmieri

President and founder, Fondazione Pensiero Solido

Mario Peserico

President, Indicam

Fabio Pompei

CEO, Deloitte Italy and Central Mediterranean

Silvio Savarese — in a personal capacity

Executive Vice President and Chief Scientist, Salesforce AI Research · Adjunct Professor of Computer Science, Stanford University · Member of the United Nations Independent International Scientific Panel on AI

Donatella Sciuto

Rector, Politecnico di Milano

Marco Tronchetti Provera

Executive Chairman, Pirelli Group

Gianmario Verona

President, Human Technopole

Giovanna Vitelli

President, Fondazione Altagamma



MilanoAI

COMITATO PROMOTORE



ASSOLOMBARDA

Conf
com



Confcommercio

Milano Lodi Monza Brianza