

Montreal Memorandum

April 2026

Preamble.

This “**Montreal Memorandum**” builds on the Digital Trust Conventions in 2024 (Paris, OECD) and 2025 (Montreal, Mila) as well as earlier discussions in a range of fora over the past few years, combining the voices of thinkers and experts from the public and private sectors, civil society and a broad range of academic and professional disciplines.

It brings together perspectives from across sectors and disciplines to articulate a shared direction for a more trustworthy, resilient, and humane digital future. In a nutshell, the Montreal Memorandum is meant to **describe and inspire a courageous, positive, even joyful digital future for shared humanity.**

Part I.

In today's digital world —

- **We observe a crisis of trust** eroded by a flood of genAI-based misleading content and deceptively human-like bots that are increasingly difficult to distinguish at scale from real people and truthful information¹. We have entered an age where the information environment is more and more distorted, designed to monetise engagement and exert influence rather than foster meaningful, cohesive and useful interaction.
- **We observe a crisis of attention** since our constant, finite human attention bandwidth meets an exponentially growing amount of content². We lack constructive and transparent mechanisms to curate the tiny slice of the information space and trust signals that we are able to process, and instead find our attention auctioned off to advertisement-centered business models.
- **We observe a crisis of markets** where vibrant competition of products and services is giving way to the dominance-preserving strategies of oligopolies and winner-takes-all narratives. We observe hype rather than a stringent and candid analysis of strengths and weaknesses, especially in new technologies such as AI.
- **We observe a crisis of common ground** driven by the polarising effects of engagement-based platform models as well as by hostile actors deliberately fostering distrust, hate and divisiveness over constructive debate. The digital world lacks resilience against persistent polarisation, while harmful behaviours scale with impunity³.

¹This could be either impossible in principle (e.g. text generated by an AI-driven chatbot does not contain sufficient data to reliably spot telltale signs of its provenance), or at least impossible in practice when the effort and cognitive load exceed what can be expected of normal users.

²Including AI-generated slop, i.e. mass-produced low value and low quality content.

³These dynamics unfold against broader social and economic conditions in which many individuals experience diminished control over their livelihoods, social status, and sense of purpose. In periods of uncertainty, the human desire to provide, to belong, and to protect those one cares about becomes more fragile. Digital systems that reward outrage, competition for attention, and performative conflict can intensify these insecurities rather than alleviate them. Recognising these shared human motivations is essential for strengthening agency and dignity.

Part II.

For tomorrow's digital world —

- **We value human agency and ambition to shape the digital world** over resigned acceptance of the status quo and technological trajectories. We see humanity as the driver, not as the bystander or overwhelmed victim, of digital transformations. Like all complex human-made systems, the digital world is inherently unstable and requires constant stewardship, governance and rebalancing.
- **We value** the health and livability of our **digital environment** for us and for future generations just like we value the health and livability of our **physical environment**⁴. Human attention is the scarce natural resource of the digital environment, while slop, deception and bots are the pollution.
- **We value human dignity and human rights** (such as equality, liberty, security, privacy, freedoms of thought, expression and communication, respect for personhood and property, health, education, and rule of law) in the **digital world** just like we value them in the **physical world**. We value accountability for upholding these rights.
- **We value** the **freedom of choosing desired degrees of digital sovereignty, self-determination and cognitive sovereignty** over being locked into unwanted dependencies. This applies at regional, national, organisational, community and individual levels, including control over our own data.
- **We value** the digital world as an inclusive **public space** rather than a fragmented collection of walled gardens; with norms, rules, governance, accountability and responsibilities that reflect its collective dimension and importance, foster democracy, and let future generations thrive.

⁴ The analogy between digital and physical environments is not merely metaphorical. Digital infrastructures, including large-scale data centres, network systems, and AI training models, carry material environmental costs in terms of energy consumption, water use, mineral extraction, and electronic waste. While this Memorandum focuses primarily on epistemic and democratic externalities, the environmental footprint of digital technologies constitutes an additional dimension of sustainability that warrants attention.

Part III.

We dare to dream!

- We dare to dream of a digital world that **brings out and preserves the best in us, both as individuals and as societies** - our ambition, our compassion, our inventiveness, our kindness, our wisdom, our collective drive for a flourishing humanity. This is the fundamental yardstick.
- We dare to dream of a digital world where **everyone's voice has an equal and real opportunity to be heard**, where human beings can stand up and be recognized with their ideas and concerns, regardless of the resources they control. This is essential for vibrant and robust democracies in the digital age.
- We dare to dream of a digital world that provides us with the infrastructure and tools to help us **judge who and what deserves our trust**.
- We dare to dream of a digital world that **encourages striving for consensus on an ever-growing set of facts** while at the same time upholding pluralism and opportunities for controversial new views to gain traction. This is essential for holding any society together, and this is also how science and engineering have progressed, based on shared access to evidence, clear terminology, and an honest understanding of what is known vs. uncertain.
- We dare to dream of a digital world **shaped by global interoperable standards and protocols rather than global monopolies** around proprietary technologies. A digital world where transparent and competitive markets drive innovation and reward human ingenuity and tenacity, and where cooperation and competition are in a healthy balance.

- We dare to dream of a digital world where technology **such as AI is a means to an end rather than an end in itself** and realises its positive potential. A world where democratic discourse and values shape the playing field rather than bowing to purported technological inevitabilities, and where AI augments human wisdom, traditions, compassion, culture and creativity rather than replacing or diminishing human expression.
- We dare to dream of a digital world that is, by its design, **structurally resilient** to the machinations of bad actors and thus does not need to suffocate constructive actors with excessive regulation.
- We dare to dream of a digital world where ownership of **capabilities and resources implies a duty to use them efficiently and for the benefit of humanity**, and to be accountable for their use. This includes physical, environmental or financial resources as well as technological edges or superior knowledge.

Part IV.

Pathways for action

This Montreal Memorandum is not a regulatory proposal or technical standard, but a shared orientation and invitation—to align action, experimentation, evaluation and governance around common principles. Having said that, it is beneficial to illustrate how such a bold, ambitious and positive digital future could be achieved through coordinated and complementary actions across three interdependent pathways: engineering systems (based on already available technologies), effective governance of their uses and impacts, and mobilising collective capacity to challenge and change entrenched incentives.

None of these pathways is sufficient on its own, and they are not a single architecture to be universally implemented. But together, they outline a direction for rebuilding trust, attention, and democratic resilience in the digital world.

These pathways are **illustrative** rather than exhaustive or authoritative. They demonstrate how the values and dreams outlined above could be put into practice. It is also important to remember that real trade-offs exist - e.g. between friction and openness, anonymity and accountability, speed and deliberation - which must be navigated transparently rather than denied.

Technological/Engineering pathways

Accountable authentic pseudonymity

We suggest the voluntary use of **accountable, authentic pseudonymity** built on existing digital identity standards, backed by trust anchors similar to existing certificate authorities, including those with distributed attestation. Such digital identities protect privacy, prove personhood and support participation integrity by limiting large-scale automated or coordinated duplication of participation, without restricting legitimate pseudonymity or role separation.

Within such an approach, everyone will make their own judgment to what extent they want to trust a person that is identified in this way, and can use open protocols to share their trust judgments with their network if they choose to. This allows reputation to spread without collapsing privacy or autonomy and can work in conjunction with provenance-based approaches.

Public knowledge graphs

We suggest building a **public knowledge graph infrastructure** as a shared epistemic backbone for the digital world. Standing on the shoulders of Wikipedia, scientific article databases and similar repositories, such non-centralised public knowledge graphs highlight the extent of agreement but also controversies and uncertainties. This infrastructure supports both human readers and AI systems, anchoring judgements in traceable sources while ensuring contestability and preserving pluralism and epistemic humility.

Democratic discourse platforms

We suggest building **democratic discourse platforms as digital commons**: non-profit, interoperable communication and civic spaces free from commercial or surveillance imperatives. Incentive structures reward constructive dialogue, compromise and consensus building to find common ground. Recommendation systems are modular and user-selectable and can be offered, e.g., by cooperatives, guilds, cultural institutions or indigenous-led initiatives that assess relevance within their domains.

Demands and proposals that garner support from a critical mass of contributors can feed into established policy making processes, e.g., being used by parliamentarians as input for drafting legislation. Thus, such platforms complement existing democratic institutions and strengthen their participatory foundations.

Resilience-by-design

We suggest establishing **resilience-by-design** as a principle for structures and processes in the digital world, creating friction for bad actors rather than relying only on hard-to-enforce regulation. Inspiration can be drawn from decades of experience in the security and safety communities.

To guide resilience-by-design, digital harm observatories monitor and analyse emerging manipulation tactics, botnets, and synthetic influence operations. Anomaly detection systems identify coordinated inauthentic behavior across networks, while incident reporting requirements ensure rapid regulatory and public response. Much of this infrastructure already exists and provides guidance and resilience indicators; we suggest strengthening it and building on top of it while maintaining transparency and proportionality.

Creator-audience-aligned content reward models

We suggest incubating a broad spectrum of models for **rewarding the creation of content that is useful for some target audience**. Such models can and should align the interests of creators and audiences, be commercially viable and sustainable, and could include, e.g., allocation mechanisms for fixed subscriptions, prosumer approaches, low-friction donations, dynamic paywalls (including for crawlers and scrapers) and, for certain types of content, matching public funding. Such models will benefit greatly from accountable authentic digital identities outlined above as a way of preventing bots and click fraud. They should avoid the need for behavioural surveillance.

We suggest addressing a broad range of content, including, e.g., news, journalistic investigations, expert advice, training materials, and research papers.

Responsible governance and collective decision making

Authority and accountability

We need legitimate and accountable processes for making **collective decisions about how digital systems and infrastructures are designed, deployed, and operated**, and who is held accountable for their impacts. Governance is not limited to regulation. It operates across multiple levels, including global rules, norms and standards, national and regional law, organisational governance, professional and sectoral standards, and community-level rule-setting. These arrangements determine who has authority to decide, how trade-offs are resolved, how differences in values and their priorities are reflected, and how accountability is enforced when harms occur.

Accountability backed by monitoring and enforcement can be implemented in a variety of ways, including liability regimes, professional standards, fiduciary duties, auditable obligations etc., and can entail e.g. responsibility for foreseeable impacts and transparency that enables public scrutiny.

Diversity of governance

Governance frameworks will necessarily **differ across jurisdictions, cultures, and political traditions**. Determining which levels of governance are appropriate for which class of decisions, and how authority and accountability are distributed across them, is one of the central challenges of the digital age. **The principles and pathways outlined in the Montreal Memorandum are intended to guide that work, not to prescribe a single model.**

Data portability and exit options

Effective governance requires meaningful **exit options as a foundation of digital freedom**, implemented in legislation rather than only on a contractual level. Without the ability to leave a platform without forfeiting one's social connections, content, and accumulated digital history, user choice is illusory. Data portability, including social graphs where appropriate, is therefore a structural prerequisite for competition, agency, and democratic plurality in the digital world.

Mobilisation and collective capacity

Governance frameworks and technical designs do not implement themselves. **Progress requires mobilisation**: the collective steps of institutions, communities, and individuals to apply pressure, build alternatives, and drive adoption of practices aligned with the values outlined in this Memorandum.

Mobilisation **takes many forms**, including coalition-building across sectors, coordinated adoption of standards and practices, strategic use of procurement and funding, professional norm-setting, and the creation of viable alternative infrastructures, starting with pilots and sandboxes for sociotechnical solutions. These efforts help shift incentives and make new arrangements practical and durable in the face of entrenched interests.

Mobilisation also includes strengthening technological literacy and AI fluency to ensure that broad segments of society can meaningfully participate in shaping and governing digital systems. This empowers individuals to know and exercise their civic rights and duties in a technology context. With this empowerment comes responsibility. In particular, freedom and fairness of expression in digital environments entail obligations toward other individuals and toward the integrity of our shared public discourse. This reciprocal balance between voice and responsibility is essential for vibrant and resilient democracies in the digital age.

Mobilisation is particularly necessary where concentrated power, market dominance, or institutional inertia resist change. Building and sustaining this collective capacity is essential for **translating governance principles and technical possibilities into lived reality**. The participation of public institutions, industry, civil society, researchers, and affected communities, each contribute distinct forms of expertise and accountability.

Initial signatories

Signatories support the overall direction and ambition even though they may not agree with every single statement.

1. **Sebastian Hallensleben**, Founder and Programme Chair, Digital Trust Convention; CEN-CENELEC, OECD
2. **Martin Hullin**, Director, Network for Technological Resilience & Sovereignty, Bertelsmann Stiftung)
3. **Pam Dixon**, Founder and Executive Director, World Privacy Forum; OECD, UN Statistics Division, WHO Health Data Collaborative
4. **Rebecca Finlay**, CEO, Partnership on AI
5. **Anna Jahn**, Executive Director, Centre for Media, Technology and Democracy
6. **Ali Hessami**, VC IEEE AIS Ethics Certification Programme and Founder, AI Community of Practice

7. **Leon Kester**, CEN-CENELEC AI Risk Management, TNO AI safety, security, meta-ethics
8. **Ojenge Winston**, Head of ACTS AI institute and founder of lab for ethics, policy and scaling of AI (LEPSAI)
9. **Talita Dias**, Senior Research Fellow, University of Exeter
10. **Gilles Fayad**, Director of initiatives, AI Commons
11. **Marta Bieńkiewicz**
12. **Axel Brando**
13. **Francesca Bosco**, Chief Strategy Officer, CyberPeace Institute
14. **Katie Evans**, Assistant Professor, Sorbonne University
15. **Jutta Juliane Meier**, Founder & CEO Identity Valley
16. **Jeannette Gorzala**, Founder & CEO ACT AI NOW
17. **Gunay Kazimzade**, AI Governance Advisor/ AI Researcher (TU Berlin)
18. **Eric Drury**, Digital Trust and Digital Identity Advisor, ForthCo.io
19. **Michael Weber**, Principal Content Quality, kununu.com
20. **Jakub Szymik**, Founder CEE Digital Democracy Watch
21. **Carjan Sabina**, AI & IP Counsel
22. **Mahsa Alimardani**, Associate Director Technology Threats and Opportunities, WITNESS
23. **Jacobo Castellanos**, Senior Coordinator Technology Threats and Opportunities, WITNESS
24. **Peter Lewis**, Director, Mindful AI Research Institute & Canada Research Chair in Trustworthy AI, Ontario Tech University
25. **David Bray**, Chair of the Accelerator and Distinguished Fellow, Stimson Center
26. **Violeta Bulc**, founder and curator, Ecocivilisation Movement

27. **David Evan Harris**, Chancellor's Public Scholar, University of California, Berkeley
28. **Romilla Syed**, Associate Professor, University of Massachusetts Boston
29. **Melisa Basol**, Sciences Po Paris; Ofcom Advisory Panel on Media Literacy
30. **Giuseppe Ugazio**, Associate Professor in Behavioral Philanthropy and Finance at Geneva Finance Research Institute
31. **Kai Zenner**, Head of Office and Digital Policy Advisor for MEP Axel Voss (European Parliament)
32. **Daniel Dobos**, Research Director, Swisscom
33. **Benjamin Prud'homme**, Vice-President, Policy, Safety and Global Affairs, Mila
34. **Sally Guy**, Director, Public Policy and Inclusion, Mila
35. **Lord Tim Clement-Jones**, Liberal Democrat spokesperson for Science Innovation and Technology House of Lords, UK