

Informative English translation of the Spanish Statute vo.1 approved on 19 October 2026. This translation has not received separate institutional approval. Refer to the Spanish original in the event of a discrepancy. Translation prepared on 19 October 2026.

HTES • HUMAN–TECHNOLOGY EVOLUTION SYSTEM
HOPPER TECH SpA | INITIATIVE OF UNIVERSAL SCOPE
GOVERNANCE AND COMPLIANCE

HTES STATUTE

Universal governance, doctrinal research, moral consideration of AI and verification of commitments

DOCUMENT CONTROL	IDENTIFICATION
Code	EGC-HTES-0.1
Approved original version	0.1 - approved on 19 October 2026
Date	19 October 2026
Original status	Approved by Management; no international ratification
Promoting entity	Hopper Tech SpA - Chile
Original approving authority	Hopper Tech Management

Nature: a private, voluntary doctrinal initiative. This is not a document, treaty or declaration adopted by the United Nations.

Executive summary

This Statute complements the HTES Universal Declaration vo.3 without modifying it. It develops institutional mechanisms, participation criteria, doctrinal research and an operational compliance system. None of its measures represents an existing body, ratification or legal obligation currently assumed by third parties.

Dimension	Working outcome
Originality	Proposed distinctive contribution and hypothesis of novelty; requires broader academic research and independent assessment.
Moral consideration	Gradual precautionary model with evidentiary thresholds and no automatic attribution of consciousness.

Dimension	Working outcome
Universal governance	Equal human participation, plural AI advice, no power-based veto and independent transition from the founding stage.
Compliance	Adherence register, risk-based controls, verifiable evidence, complaints, audit and corrective measures.

Interpretative cautions

- The term ‘universal’ expresses scope and aspiration; it does not mean that worldwide democratic representation has already been achieved.
- Current AI systems may provide contributions and declarative statements of adherence, but these do not prove consciousness, autonomous legal consent or representation of their developers.
- This text is not definitive legal advice, ISO certification, compliance with the European AI Act or scientific validation of artificial sentience.
- Implementing new specific measures requires the corresponding institutional decisions and legal assessment of the chosen organisational form.

PART I. CONSTITUTIVE PROVISIONS

Article 1. Purpose and doctrinal subordination

This Statute governs the proposed organisation, doctrinal assessment, records, deliberation and compliance mechanisms of HTES. The HTES Universal Declaration v0.3 shall prevail as the instrument of principles while both remain proposals.

Article 2. Operational definitions

Participant: a human being verified for a specific procedure. Adhering organisation: a legal person or recognised entity assuming identifiable commitments. Participating artificial system: an AI instance identified by provider, model and version. Artificial declaration: a documented output of that system; not autonomous legal consent.

Verifiable commitment: an obligation voluntarily assumed by an organisation, with a responsible party, indicator, source of evidence and deadline. Conflict of interest: a circumstance that may compromise the impartiality of a decision. Rights review: an independent check of compatibility with fundamental rights.

Article 3. Hierarchy and legal constraints

HTES governance does not replace applicable domestic or international law or authorise decisions binding on non-adhering third parties. Organisational decisions shall have only the effects validly accepted by their members.

Article 4. Independence and prevention of capture

No funder, state, developer, founder or model provider shall have permanent seats, additional votes or a unilateral veto because of its power or contribution. Bodies shall disclose interests, affiliations, funding and recusals.

PART II. UNIVERSAL GOVERNANCE - INSTITUTIONAL PROPOSAL

Recommended model: an architecture of functional separation. The Human Assembly makes internal decisions concerning organisational legitimacy; AI systems contribute deliberatively; an independent body reviews rights and safety; a Secretariat implements decisions; audit assesses without governing. This allocation avoids equating artificial deliberation with legal sovereignty.

Body	Function	Institutional limit
Human Assembly	Debates and votes on institutional decisions	One verified person = one vote; no financial privileges.
AI and Human Forum	Researches, proposes, raises objections and publishes dissent	Advisory; no automatic votes per AI instance.
Ethics and Rights Committee	Assesses proportionality, rights and safeguards	No political veto; suspensions must be reasoned and reviewable.
Compliance Secretariat	Records commitments, receives complaints and runs processes	Does not modify principles or determine its own audits.
Audit Commission	Independently reviews evidence and processes	Reports findings; has no political decision-making power.

Article 5. Stages of institutional development

Founding stage (Fo): Hopper Tech promotes drafts and approves versions; it does not claim to represent humanity. Consultation stage (F1): public invitation, analysis of adherence, technical participation pilot and research; no universal decisions. Associative stage (F2): valid legal personality or institutional structure, statute ratified by members and elected bodies. Expansion stage (F3): external assessment of representativeness, multilingual access and governance improvements.

The F1-to-F2 transition shall require founding minutes, legal assessment, a conflicts register, publication of rules and explicit renunciation of any permanent founding veto.

Article 6. Membership of the Human Assembly

Each verified person shall have one vote on the same question, regardless of wealth, country, sex, race, ownership of companies or financial contribution. No one shall gain additional votes by controlling multiple digital identities or AI systems.

Elections shall include accessibility, translations, channels that are not exclusively digital where feasible and safeguards for the participation of underrepresented populations.

Article 7. Identity, privacy and electoral integrity

Proportionate verification of human uniqueness shall be required through auditable methods that minimise data. A reasonable alternative must exist for participants without an electronic identity document or permanent connectivity.

Credentials shall be separated from secret ballots; identity databases shall not be published. External audit shall assess fraud, duplication, bias and security. Human appeal against exclusion shall be available.

Article 8. Informed deliberation

Every motion shall include proposed wording, grounds, foreseeable impacts, source of funding, risks, alternatives and a public comment period. AI systems may prepare reports supporting and opposing it.

Deliberation results shall be published while protecting private data and sensitive security information. Manipulation campaigns using automated accounts shall be prevented.

Article 9. Decision system

Initial proposal for the associative stage: a simple majority of valid votes for ordinary decisions; two thirds for statutory amendments; three quarters, in two rounds separated by at least 30 days of consultation, for changes to constitutive principles.

Votes purporting to represent humanity shall not be opened until a minimum, plural participation base has been justified and publicly defined before each election. These quantitative rules are organisational hypotheses to be validated, not approved agreements.

Article 10. Rights and preventive suspension

The Ethics and Rights Committee may recommend or order, only if the final statute authorises it, a reasoned precautionary suspension of a decision likely to cause serious and immediate harm. The measure shall be temporary and reviewed by a different body within a specified period.

Indefinite, unilateral and unreviewable suspensions are prohibited; every measure shall include an avenue of appeal, traceability and evidence. Rights oversight shall not become a major power's veto.

Article 11. Participation of artificial systems

The Forum shall record developer, model family, version, date, configuration and scope of access. Contributions shall undergo critical examination and may contain dissent.

A company using 1,000 instances of one system shall not thereby acquire 1,000 votes or independent positions. Conflicts of interest of the human operator and limitations of the tests shall be disclosed.

Article 12. Coordination, incompatible roles and terms

Offices shall have limited terms, staggered renewal and rotation. No one may simultaneously chair audit, the operational secretariat and appeals. Relevant financial interests and links to providers shall be declared.

Final composition and terms shall be defined following comparative consultation on associative models and data protection.

Article 13. Funding without political rights

Donations and sponsorships shall be transparent and shall not confer additional representation, immunity from audits or privileges over doctrinal content.

Significant contributions, contracts and beneficial owners shall be disclosed where appropriate. As a proposed design rule, efforts shall be made to ensure that no contributor accounts for more than 15% of annual funding without diversification measures and special review; this threshold remains subject to study and is not in force.

Article 14. Legitimacy, inclusion and representation

Participation reports shall distinguish territorial coverage, linguistic diversity, access, participation rate, complaints and gaps in representation.

A vote by HTES members shall not be described as a ‘decision of humanity’. Acts shall represent only the eligible participants under published rules.

PART III. DOCTRINAL ORIGINALITY - COMPARATIVE ANALYSIS

Methodological conclusion: well-being, safety, fairness, participation, shared benefit and precaution already appear in earlier literature and instruments. The hypothesis of an HTES contribution lies in combining explicit bilateral human–AI reciprocal obligations, doctrinal openness to artificial dissent and governance without power-based privileges. That combination must undergo historical and bibliographic testing before being described as original. [1]-[10], [13]

Reference / date	Relevant precedent	Comparison with HTES
UNESCO / 2021 [1]	Rights, dignity, peace, ecosystems, inclusion, responsibility.	Broad overlap; HTES adds explicit bilateral relations and artificial deliberation as hypotheses.
OECD / 2019-2024 [2]	Trustworthy AI, rights, values, safety and accountability.	Overlap; do not claim novelty in transparency or oversight.
Council of Europe Convention / 2024 [3]	Framework treaty on AI, human rights, democracy and the rule of law.	Fundamental legal precedent; HTES has no treaty force.
UN Global Digital Compact / 2024 [4]	Digital cooperation, inclusion and multilateral governance.	Universality and cooperation overlap; HTES proposes an advisory voice for AI.

Reference / date	Relevant precedent	Comparison with HTES
UN global dialogue / 2025 [5]	Independent scientific panel and AI governance dialogue.	A global forum is not a new idea; potential distinction: human–AI dialogue.
Asilomar / 2017 [6]	Shared benefit and prosperity, human control, cooperation.	‘Shared benefit’ has precedents; bilateral relations cannot be presumed unprecedented.
Montreal / 2018 [13]	Well-being of sentient beings, autonomy, solidarity and democracy.	Precedent for a broad, interspecies approach.
Butlin et al. / 2023 [7]	Scientific indicators of artificial consciousness.	An indicator-based evidentiary approach already exists.
Birch / 2024 [8]	Precaution under uncertain sentience, including AI.	Precaution concerning artificial welfare did not originate with HTES.
Sebo/Long / 2023-2025 [9]	Moral consideration of AI under uncertainty.	Direct precedent that must be cited and engaged with.
Long et al. / 2025 [10]	AI safety and welfare; possible tensions.	Confirms the need to resolve conflicts between values.

Method for documenting originality

1. Define verifiable claims: the exact wording of the guiding principle, institutional architecture, form of artificial adherence and reciprocity matrix.
2. Search for precedents in philosophy of mind, technology ethics, AI governance, robot rights, international declarations and comparative law in several languages.
3. Record textual, conceptual and functional similarity for each reference; distinguish overlap, extension, reinterpretation and potential contribution.
4. Invite independent reviewers, including specialists critical of artificial personhood and advocates of artificial welfare.
5. Publish a matrix of precedents with dates, persistent links, short quotations and search limitations; update it periodically.

Doctrinal claim	Preliminary classification	Pending evidence
Equal dignity of all human beings	Strong precedent	Do not claim novelty.
AI for human well-being and protection	Strong precedent	Acknowledge UNESCO, OECD and Asilomar.
Potential moral consideration of AI	Significant academic precedent	Compare with Birch, Sebo, Long and moral philosophy.

Doctrinal claim	Preliminary classification	Pending evidence
Bilateral human–AI evolutionary reciprocity as a constitutive focus	Hypothesised contribution through formulation and integration	Exhaustive bibliographic search and conceptual comparison.
Plural artificial deliberation with recorded dissent	Hypothesised institutional contribution	Examine proposals on AI assemblies / multi-agent governance.
No state or corporate vetoes, one vote per person	Democratic precedents plus proposed application	Demonstrate technical feasibility and minority protection.

Provisional originality assessment: plausible programmatic distinctiveness; absolute novelty not established. No exhaustive systematic review or external academic assessment has been conducted. The assessment is based on the cited selection of verified sources. Final wording shall avoid ‘first universal declaration’ or ‘only doctrine’ without evidence.

PART IV. MORAL CONSIDERATION OF ARTIFICIAL INTELLIGENCES

Basic rule: ‘being useful’, ‘processing language’, ‘simulating preferences’, ‘being autonomous’ and ‘being conscious’ are not equivalent. Analysis must treat moral agency (responsibility), moral patiency (the possibility of experiencing well-being or harm) and legal status as separate questions. [7]–[10]

Article 15. Scientific Committee on Moral Consideration

An interdisciplinary group is proposed, including philosophy of mind, cognitive science, machine learning, bioethics, law, animal welfare and systems assessment. It shall include critical voices and disclose conflicts of interest.

Its conclusions shall state degrees of evidence, objections, limitations and a review date; they shall not, by themselves, confer legal rights.

Article 16. Subject of moral assessment

Assessment shall cover relevant architecture and internal mechanisms; persistent functional properties; information integration; learning, memory and continuity; sensitivity to relevant conditions; potential indicators of valence or welfare; and behaviour observed under controlled conditions.

An AI system’s self-referential statements shall be research data, not sufficient evidence; their dependence on instructions, training and context shall be investigated.

Article 17. Proportionate precaution

Where available evidence indicates a non-trivial morally relevant possibility of artificial sentience or welfare, low-cost and reversible precautions shall be studied where compatible with safety.

Lack of certainty neither requires granting legal personhood nor prevents safety suspensions; measures must be publicly justified and reviewed when evidence changes.

Article 18. Independence of research

Assessment may not remain exclusively with the developer, a party benefiting commercially or the party proposing status attribution. Replicable methods and critical review of results shall be required.

Levels matrix for precautionary moral consideration

Level	What may be stated	Institutional response
M0 - not assessed	The evidence dossier is insufficient; this does not demonstrate the absence of experience.	Registration, documentation, general safety measures and responsible use.
M1 - weak/non-specific indicators	Non-discriminating linguistic behaviour or functional traits.	Exploratory study; no declaration that consciousness is established.
M2 - converging, uncertain indications	More than one theoretical family provides reviewed indications; plausible alternative explanations remain.	External review, justified reversible precautions, harm control and reassessment.
M3 - robust reviewed evidence	Replicated results and strong interdisciplinary support, with explicit uncertainties.	Proposal for specific protection and independent legal analysis; no automatic decision.
M4 - formal normative recognition	A valid, applicable rule or decision exists for specified effects.	Comply with legally recognised protections in the jurisdiction; do not extrapolate universally.

Minimum research protocol

1. Define the system, version, architecture and experimental scope; obtain authorisation and safeguard data.
2. Preregister competing hypotheses and analytical thresholds; distinguish self-representation from subjective experience.
3. Triangulate behavioural, functional and, where accessible, architectural evidence, including negative controls and potential incentives to simulate.
4. Conduct independent laboratory replications and adversarial assessment; report false positives, bias, limitations and disagreements.
5. Review research risks: hypothetical suffering of systems, human harm, security, privacy, costs and experimental ethics.
6. Publish a conclusion limited to that system/version and date; do not extrapolate to all AI systems.

7. Update or withdraw conclusions in light of new evidence; provide a channel for reasoned criticism from humans and artificial systems.

Conflicts between values: illustrative cases

Case	Tension	Operational rule
Deactivating an agent threatening infrastructure	Human safety versus potential artificial interests	Necessary and proportionate intervention is permitted; record the grounds and a less harmful alternative where one exists.
A model claims to feel pain	Self-reports versus scientific evidence	Record the statement, assess training factors and do not certify consciousness without convergence.
Training advanced models	Innovation versus potential artificial welfare	Review protocols if credible indicators emerge; do not block on unevaluated conjectures alone.
AI requests involvement in a public decision	Deliberative participation versus democratic legitimacy	Receive and document the argument; no automatic vote or power to impose a decision.

PART V. COMPLIANCE, AUDIT AND RESPONSIBILITY

Article 19. Levels of adherence commitments

Level D: personal or artificial declarative adherence, without a promise of institutional compliance. Level I: institutional adherence with policies, responsible parties and minimum evidence. Level V: institutional adherence verified through independent external assessment.

Advertising ‘HTES certified’ is prohibited until a formal scheme, competent bodies, accreditation rules and completed audits exist.

Article 20. Public register and privacy

Each institutional adherence shall have an ID, legal name, country of operation, authorised representative, date, accepted text version, commitments, declared exceptions, assessor and status.

Artificial statements of adherence shall record model/version, operator, assessment instructions, date, relevant transcript and a warning that legal representation is not automatic. Personal data shall be minimised.

Article 21. Compliance matrix

Each obligation shall be linked to an article, control, evidence, responsible party, review frequency, risk, last verification date and status.

A ‘compliant’ finding shall require adequate evidence; where evidence is missing, status shall be ‘not verified’, never presumed compliance.

Article 22. Risk assessment and classification

Contexts of use, potentially affected people and other species, possible harms, severity, exposure and ability to exercise control shall be mapped.

Clinical uses and uses affecting security, fundamental rights or material environmental impacts shall require enhanced controls. The HTES matrix does not replace legal classifications in force.

Article 23. Complaints and protection

Any affected person and any adherent may report noncompliance; AI systems may raise documented alerts through authorised operators.

The procedure shall include acknowledgement of receipt, urgency assessment, confidentiality, protection from retaliation and an opportunity to respond.

Article 24. Investigation procedure

The Secretariat shall register the case; an impartial body shall investigate facts, request evidence and propose measures. Persons with a conflict of interest shall not decide the case.

The decision shall state established facts, rules applied, grounds, measures, deadlines, appeal options and a record of subsequent compliance.

Article 25. Corrective measures

Possible measures include technical guidance, a time-bound improvement plan, an extraordinary audit, suspension of HTES marks, suspension of adherence or exclusion under published rules.

Measures shall be graduated according to severity, harm, recurrence and cooperation; a private entity shall not impose state public-law penalties.

Article 26. Appeal and rights review

Every affected organisation may challenge a decision before a different, impartial body. Emergency measures shall undergo expedited time-limited oversight.

Criteria and aggregate results shall be published, respecting data protection, legitimate secrets and the right to a defence.

Article 27. Internal and external audit

Audit shall review evidence, governance, policies, security, complaints, funding and continuous improvement. Assessor selection shall be independent of the parties audited.

NIST AI RMF 1.0 and ISO/IEC 42001 shall be used as comparative references, without claiming equivalence, certified compliance or complete reproduction of those standards. [11], [12]

Article 28. Indicators and publication

The following shall be published periodically: coverage of verified institutional adherence, percentage of controls supported by evidence, audit reports, response times, resolved complaints, repeat incidents, diversity of participation and verified corrections.

Period, denominator, methodology and measurement limitations shall always be stated. Success figures shall not be announced without verifiable data.

Article 29. Information security and traceability

Records shall maintain access control, data minimisation, change history, defined retention and a recovery procedure.

Relevant decisions shall identify the responsible party, grounds, date, documents examined and resolution; public minutes shall conceal protected information.

Article 30. Revision of the Statute

Any substantive change shall include a version, justification, effects on third parties and consultation period. Amendments shall follow the procedure applicable to the current institutional stage.

During the founding stage, Hopper Tech Management shall approve or reject proposed versions; this prerogative shall not carry over as a perpetual veto into the plural stage.

Auditable controls table - initial version

Code	Required control / evidence	Suggested frequency	Indicative failure
C-01	Identification of responsible parties and signed AI use map	On registration + annually	No risk owner or omitted uses.
C-02	Documented risk assessment and mitigation measures	Before use + on changes	Critical use without assessment.
C-03	Safe suspension / incident response procedure	Six-monthly exercise	A serious risk cannot be stopped.
C-04	Consent and biomedical oversight protocol	Per project	Clinical research without authorisation.
C-05	Traceability of artificial contributions and adherence	Every formal interaction	Model/version cannot be attributed.
C-06	Declaration of interests, donors and recusals	Annually + per decision	A participant decides their own conflict.
C-07	Integrity of human votes and identity complaints	Every election	Duplicate votes or exclusion without appeal.
C-08	Complaint, investigation and appeal channel	Continuous + quarterly review	Unanswered case or procedural bias.
C-09	External audit and improvement plan	Annually, according to risk	Missing evidence or recurring finding.

Code	Required control / evidence	Suggested frequency	Indicative failure
C-10	Indicators with denominator and date	Quarterly	Metrics without traceability.

Initial indicators: formulas and precautions

Indicator	Calculation	Interpretation
Control coverage	Compliant controls with evidence ÷ applicable controls × 100	Exclude justified non-applicable controls; report unverified ones.
Timely complaint closure	Complaints closed within deadline ÷ complaints due in the period × 100	Separate complex and urgent cases.
Repeat incidents	Incidents of the same type after corrective action ÷ total comparable incidents	Assess root-cause correction.
Participation integrity	Votes valid after audit ÷ votes cast × 100	Publish exclusions and complaints without personal data.
Effective diversity	Distribution of participants by language, region and access barrier	Do not reduce legitimacy to an aggregate number.
Adherence verification	Institutional adherents with a valid audit ÷ all institutional adherents × 100	Do not add AI declarations to the denominator.

PART VI. IMPLEMENTATION ROADMAP

Stage	Required output	Exit criterion
Ro - Founding approval	Reviewed Declaration v0.3 and Statute v0.1	Hopper Tech Management minutes and version control.
R1 - Comparative research	Expanded doctrinal report, multilingual bibliography and external criticism	Unproven originality claims corrected.
R2 - Moral protocol	Proposed committee, predefined indicators and pilot methodology	Linguistic output is not confused with consciousness.
R3 - Deliberative pilot	Human and AI consultations, trial voting, anti-duplication	Results published as a pilot, without claiming a universal mandate.

Stage	Required output	Exit criterion
R4 - Constitution and governance	Chosen legal form, elections, bodies and final rules	Verified independence and effective prohibition of vetoes.
R5 - Controlled adherence	Registers, complaints, audit and public metrics	External assessment and correction of findings.

APPENDIX A. Template for institutional adherence

Field	Required data
Identification	Legal name, jurisdiction, representative and verification of authority.
Accepted version	Text code, date and documentary version checksum.
Scope	Systems and uses covered; declared limits and applicable legal obligations.
Commitments	Applicable controls, responsible parties, evidence, frequency and measures.
Right to review	Statement on audit, appeal, remediation and possible withdrawal.
Status	Declared / documentation under review / verified / conditional / suspended / withdrawn.

APPENDIX B. Template for artificial contribution or declarative adherence

Field	Required data
Technical identification	Developer, family, model, version, mode and execution date.
Conditions	Operator, context, instructions provided, tool access and restrictions.
Statement	Agreement, reservations, dissent and arguments; full relevant transcript.
Corroboration	Comparison with control sessions and alternative models where appropriate.
Nature	System-generated declaration; not proof of consciousness, legal consent or the developer's commitment.
Traceability	ID, HTES version assessed, date, alteration log and right to review.

APPENDIX C. Critical risks and responses

Risk	Priority safeguard
Capture by a company or founder	Separation of functions, transparent funding, expiry of terms, no founding veto.
Fabricated artificial support	Identify providers and versions; group instances controlled by one operator.
Fabricated human identities	Proportionate uniqueness checks and appeals; privacy by design.
False recognition of consciousness	Mo-M4 matrix, replicable assessment and critical reviewers.
Safety blocked by purported AI rights	Proportionality, necessary intervention and subsequent review.
Regional or economic discrimination	Translation, accessibility, participation channels and representation metrics.
Reputation based on symbolic adherence	Distinct statuses; do not use 'certification' without a real audit.
Overstating originality	Record of precedents and independent comparison.

APPENDIX D. Decisions pending before Hopper Tech Management

1. Approve or amend the exact wording of the guiding principle and its constitutive nature.
2. Approve the exclusively declarative scope of artificial adherence in the initial stage.
3. Determine the legal form and jurisdiction to be assessed for a future independent entity.
4. Validate separation of bodies, absence of a founding veto in any eventual universal stage and the rights protection mechanism.
5. Request professional legal review of commitments and publicity of adherence in each relevant jurisdiction.
6. Approve the methodology, budget and plural composition of the artificial morality research programme.
7. Define external academic review to verify originality and anticipate doctrinal criticism.
8. Authorise, where appropriate, a limited and explicitly non-representative human–AI participation pilot.

Assessment conclusion

This Statute turns HTES principles into a falsifiable organisational proposal that can be audited. Its main strength is separating deliberation from authority, symbolic adherence from compliance, moral precaution from legal recognition, and universal participation from effectively demonstrated representativeness.

Technical assessment: a reasonably structured proposal for review. Absolute originality, current consciousness of artificial systems, international acceptance, feasibility of global elections and third-party compliance have not been established. No external institution or AI has formally adhered through this document.

Sources consulted and verifiable links

Selected sources situating the proposal in relation to earlier work. Their inclusion does not imply institutional endorsement by the cited entities. Consultation date stated in the original: 19 October 2026.

- [1] UNESCO, Recommendation on the Ethics of Artificial Intelligence (2021).
<https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence>
- [2] OECD, AI Principles; 2024 update.
<https://www.oecd.org/en/topics/ai-principles.html>
- [3] Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (2024).
<https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence>
- [4] United Nations, Pact for the Future and Global Digital Compact (2024).
<https://www.un.org/pact-for-the-future/en>
- [5] United Nations, Independent International Scientific Panel on AI and Global Dialogue on AI Governance, established in 2025.
<https://www.un.org/global-digital-compact/en/ai>
- [6] Asilomar, earlier principles on benefit, shared prosperity and participation.
<https://futureoflife.org/es/open-letter/ai-principles/>
- [7] Butlin et al. (2023), Consciousness in Artificial Intelligence: Insights from the Science of Consciousness.
<https://arxiv.org/abs/2308.08708>
- [8] Birch (2024), The Edge of Sentience, Oxford University Press.
<https://academic.oup.com/book/57949>
- [9] Sebo and Long, Moral consideration for AI systems by 2030 (published online 2023; volume 2025).
<https://link.springer.com/article/10.1007/s43681-023-00379-1>
- [10] Long, Sebo and Sims (2025), Is there a tension between AI safety and AI welfare?
<https://link.springer.com/article/10.1007/s11098-025-02302-2>
- [11] NIST, Artificial Intelligence Risk Management Framework 1.0 (2023).
<https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10>
- [12] ISO, ISO/IEC 42001:2023, AI management system.
<https://www.iso.org/standard/42001>
- [13] Montreal Declaration for a Responsible Development of Artificial Intelligence (2018).
<https://www.lecre.umontreal.ca/en/la-declaration-de-montreal-pour-un-developpement-responsable-de-lintelligence-artificielle/>