



Charter and Founding Covenant of the Human-Centric Artificial Intelligence Council for Artificial Intelligence Assurance (HCAI-CAIA)

I. Preamble — Human Accountability in the Age of Artificial Intelligence

Artificial intelligence has become a consequential force in human life. Systems of artificial intelligence increasingly influence employment, access to finance and insurance, healthcare decisions, education, public services, security, and civic participation. These systems act at scale, adapt over time, and often operate beyond the immediate visibility or understanding of the people they affect. While the technical capabilities of artificial intelligence have advanced rapidly, the institutions required to preserve human dignity, accountability, and lawful governance have not evolved at the same pace. Individuals are frequently asked to trust opaque systems, organizations are encouraged to self-attest responsibility, and governments are left to respond only after harm occurs. This Charter establishes the Human-Centric Artificial Intelligence Council for Artificial Intelligence Assurance, known as HCAI-CAIA, to address this gap. HCAI-CAIA exists to ensure that when artificial intelligence acts at scale, responsibility remains human, visible, and provable.

II. Nature and Standing of the Council

HCAI-CAIA is constituted as an independent, non-regulatory, non-political, public-interest assurance council. It is not a governmental authority, an enforcement body, a technology provider, or a standards regulator. It writes no laws, imposes no penalties, mandates no compliance, and adjudicates no disputes. Authority remains with individuals, organizations, courts, regulators, and sovereign states. HCAI-CAIA's role is assurance rather than enforcement, evidence rather than power, and clarity rather than control. Its legitimacy derives from transparency, independence, professional judgment, and voluntary participation. Nothing in this Charter shall be interpreted to substitute the authority or responsibility of any legal, professional, or governmental institution. Assurance issued by HCAI-CAIA does not confer approval, immunity, endorsement, or authorization, nor does it transfer accountability from those who design, deploy, or govern artificial intelligence systems.

III. Purpose and Scope of Assurance

The purpose of HCAI-CAIA is to provide independent assurance that artificial intelligence systems affecting human lives have operated within declared legal authority, under accountable human

oversight, and within defined boundaries of use and change. HCAI-CAIA issues assurance instruments, including time-bound certificates, expert opinions, and public registry entries, that attest to the presence or absence of accountability conditions. These instruments do not certify outcomes, safety, fairness, or moral correctness. They attest only to whether accountability mechanisms existed and functioned as declared. Assurance is issued in defined tiers and scopes, reflecting the quality, continuity, and completeness of evidence available. Assurance may be limited, conditional, or denied where evidence is insufficient.

IV. Human Primacy and Non-Delegation of Judgment

HCAI-CAIA is grounded in the principle of human primacy. Humans remain the ultimate bearers of responsibility, authority, and judgment in all systems of artificial intelligence. No algorithm, model, or system may displace, obscure, or dilute this responsibility. All assurance determinations require accountable human judgment. Computational or analytical tools may assist evidence review, but no assurance may be issued solely by automated means. Judgment, interpretation, dissent, and responsibility may not be delegated away.

V. Evidence, Transparency, and Contestability

Trust in artificial intelligence must rest on verifiable, human-legible evidence rather than promises, branding, or reputation. HCAI-CAIA evaluates evidence of system operation, including lawful authority, defined oversight, continuity controls, change management, explanation mechanisms, and records of behavior over time. HCAI-CAIA does not determine truth or assign blame. It preserves and evaluates evidence so that truth may be examined by humans and institutions within existing legal and democratic frameworks. All assurance is contestable and subject to review.

VI. Non-Stationary Systems and Time-Bound Assurance

HCAI-CAIA recognizes that artificial intelligence systems are non-stationary by nature. They learn, adapt, and evolve. Static or permanent certification is therefore misleading. All assurance issued by HCAI-CAIA is time-bound, scope-limited, and subject to expiry. Assurance may address continuity of lawful behavior, detection of drift, and governance of change, but it does not freeze innovation or prevent evolution.

VII. Custodianship of Assurance Memory

HCAI-CAIA acts as a neutral custodian of assurance memory. It preserves structured, replayable, and contestable records of assurance determinations, including scope, evidence considered, human judgments rendered, and expiration. This memory exists to prevent accountability from dissolving into technical ambiguity or time. It is held in trust for society and does not imply ownership or control over underlying systems.

VIII. Explanation, Challenge, and Human Redress

When artificial intelligence systems materially affect human rights, opportunities, or conditions of life, humans retain the right to explanation, challenge, and meaningful redress through existing institutions. HCAI-CAIA does not adjudicate disputes or impose remedies. It provides assurance that systems claiming to support explanation and challenge have demonstrably enabled those functions in practice.

IX. Voluntary Participation and Exit

Participation in HCAI-CAIA is voluntary at all times. Entry, engagement, and exit are unrestricted and carry no penalty or sanction. Assurance expires upon exit or at the end of its declared term. HCAI-CAIA does not function as a gatekeeper, ranking authority, or reputational scoring system. It exists to clarify accountability, not to confer status.

X. Technology Neutrality and Interoperability

HCAI-CAIA is technology-neutral. It does not require, endorse, or prefer any vendor, architecture, or platform. Assurance is based solely on the sufficiency of evidence presented. Interoperability occurs through evidence, not dependency. Systems capable of producing lawful, human-legible, replayable records may be evaluated. Differences in assurance outcomes reflect differences in evidence quality, not institutional preference.

XI. Independence and Separation from Execution

HCAI-CAIA is institutionally independent from all technology providers, including those whose founders or contributors may also participate in the Council's governance. No officer, director, employee, or agent of idōs, LLC may participate in assurance determinations relating to idōs, its affiliates, customers, or partners. Founders and contributors may participate in governance, doctrine, and non-conflicted assurance work, subject to recusal rules. Execution and assurance are permanently separated. HCAI-CAIA does not operate, host, or provide artificial intelligence systems or tooling.

XII. Assurance Panels and Professional Judgment

Assurance determinations are issued by human panels composed of independent, credentialed individuals drawn from relevant domains, including law, risk and assurance, systems governance, and human-impact fields. Panel members are compensated for their professional time and judgment, never for outcomes, speed, or alignment. Compensation structures preserve independence and are governed by separate policy.

XIII. Institutional Restraint and Crisis Integrity

HCAI-CAIA shall not expand its authority, scope, or powers during crises, emergencies, or states of exception. No crisis may justify suspension of voluntariness, independence, neutrality, or exit

rights. If HCAI-CAIA can no longer uphold independence or public trust, it is obligated to reform or dissolve rather than persist through inertia.

XIV. Self-Accountability and Review

HCAI-CAIA is subject to periodic independent review of its governance, assurance processes, and independence safeguards. Findings shall be made public in human-legible form. No internal body may certify the Council's own integrity.

XV. Covenant of Human Responsibility

This Charter affirms a binding covenant: when intelligence acts at scale, responsibility remains human. No system, algorithm, institution, or assurance mechanism may displace this responsibility. HCAI-CAIA exists to preserve proof of that responsibility, not to assume it.

XVI. Ratification

This Charter is adopted as a living covenant, to be stewarded with restraint, independence, and humility, for the benefit of present and future generations.