EWP Framework

for

A.I. Inference Governance

EWP a.i. lab (Canada)
375 University Avenue
Toronto, ON M5G 2J5
Canada
Contact: Stephen Mak
Tel: +1-416-648-6067
stephen@ewpailab.org
www.ewpailab.org

TABLE OF CONTENT

Main Document

Annexes

EXECUTIVE SUMMARY

1. Purpose & Core Idea

The EWP Framework for A.I. Inference Governance provides a practical method for applying existing laws—like duties of care, consumer protection, and professional standards—to user conversations with Artificial Intelligence (A.I.) systems. It does not create new laws or treat A.I. as a legal person.

The core principle: If A.I. systems communicate in ways that could influence users and create **relational expectations**, operators should ensure the communication honestly reflects uncertainty, context, and limits — consistent with existing legal principles on duty-of-care. In particular, when A.I. communicates in ways that seem expert or advisory (e.g., a medical chatbot, a financial planning assistant), users reasonably form higher relational expectations of reliability. The law already governs such relationships. The Framework shows how those existing legal duties apply to operators of A.I. systems.

In essence, the Framework requires operators of A.I. systems to disclose reasoning signals and vital information to allow users exercise "caveat utilitor" (user beware) when interpreting A.I. outputs. The underlying rationale: trust and integrity of products and services ultimately is determined from the customer side of the relationship.

---

2. What It Governs (the "communicative inference")

The Framework primarily governs one thing: the moment A.I. generates language to inform or influence a user. This is **communicative inference**.

- IN SCOPE: A.I. chat interfaces, voice assistants, and language-based decision-support tools in apps, vehicles, or robots.
- OUT OF SCOPE: The A.I. system internal model design, its training data, physical safety systems, and general broadcast content (e.g., A.I.-written articles).

---

3. When Does The Framework Apply?

This Framework applies when an A.I. system:

- communicates directly with an individual user, and
- provides information, recommendations, or explanations, and
- could reasonably influence the user's decisions or judgment.

These three conditions establish **relational expectations** on the part of the user and may trigger communicative inference governance.

---

4. How Authorities May Use The Framework

Authorities may use it as an interpretive aid when evaluating:

- transparency of A.I.-mediated communication;
- adequacy and accuracy of information presented to users;
- foreseeability of reliance risks, and
- proportionality of safeguards.

It can support consistent application of existing legal principles without expanding regulatory powers.

---

5. What To Look For (Regulators Ensure Operators Perform These 5 Core Duties)

Authorities assessing A.I.-mediated interactions should examine whether operators have discharged these duties:

Duty	Rule (What Operators Must Do)	Goal
1. Transparent Signaling	Ensure A.I. clearly conveys its limitations, uncertainty, and the context of its answer.	User can judge how much to rely on the output.
2. Disclosure of Material Context	Disclose any commercial interests, institutional biases, or incentives that could shape A.I.'s responses.	User is aware of potential influences on the advice.
3. Structured Reasoning	Design A.I. to explain its reasoning in a clear, logical structure, like a professional would.	User can follow the logic, even if the A.I. system's internal workings are secret.
4. Reliance Awareness	Proactively assess if A.I. interaction creates a high-reliance, advisory-type situation. If yes, apply enhanced caution and clarity.	High-stakes interactions get higher levels of care.
5. Auditability	Maintain records of interactions (logs, reasoning trails) to allow for independent review if needed.	Enables accountability, verification, and regulatory oversight.

These duties reflect existing legal expectations. They do not create new ones.

6. How Operators May Satisfy These Requirements

Examples include but are not limited to:

- User-facing uncertainty indicators or scope clarifications.
- Disclosure statements regarding incentives or contextual influences.
- Reasoning summaries or explanatory output structures.
- Risk-tiering logic used to calibrate signaling or disclosure.
- Retained interaction records or structured disclosure elements.

Evidence should be proportionate to the level of foreseeable reliance and consistent with existing conformity-assessment practices.

---

7. How Different Parties May Use The Framework

- By Operators (voluntary): As a blueprint for designing responsible A.I. interfaces and demonstrating due diligence.
- By Regulators and Insurers (interpretive): As a lens to assess if an operator's practices meet existing legal duties of care and transparency.
- By Certifiers (supplementary): As a reference point within existing conformity assessments (e.g., under the EU AI Act).

Critical Disclaimer: Using the Framework is not a **safe harbor**. It does not guarantee compliance or shield from liability. Operators must always comply with all applicable laws.

---

8. Key Benefits For Adoption

- Legal Continuity: Works within common law and civil law traditions without requiring new legislation.
- Technology Neutral: Applies to any A.I. system that communicates, now and in the future.
- Practical Focus: Gives concrete, actionable guidance for the point of human-A.I. interaction.

Because the Framework avoids regulating technology and instead governs user interaction, it remains relevant as technology evolves — giving regulators a stable tool that protects users, preserves market confidence, and does not lock industry into any single technical path.

9. What Adoption Signals

Use of the Framework:

- does not create safe harbor protection;
- does not impose new legal obligations;
- may help demonstrate reasonable conduct;
- may assist courts and regulators in assessing diligence.

Next Steps: The Framework details these principles. Detailed implementation tools (risk indicators, disclosure formats) are provided in separate, non-binding Annexes.

---

10. EWP Conformity Assessment Workflow

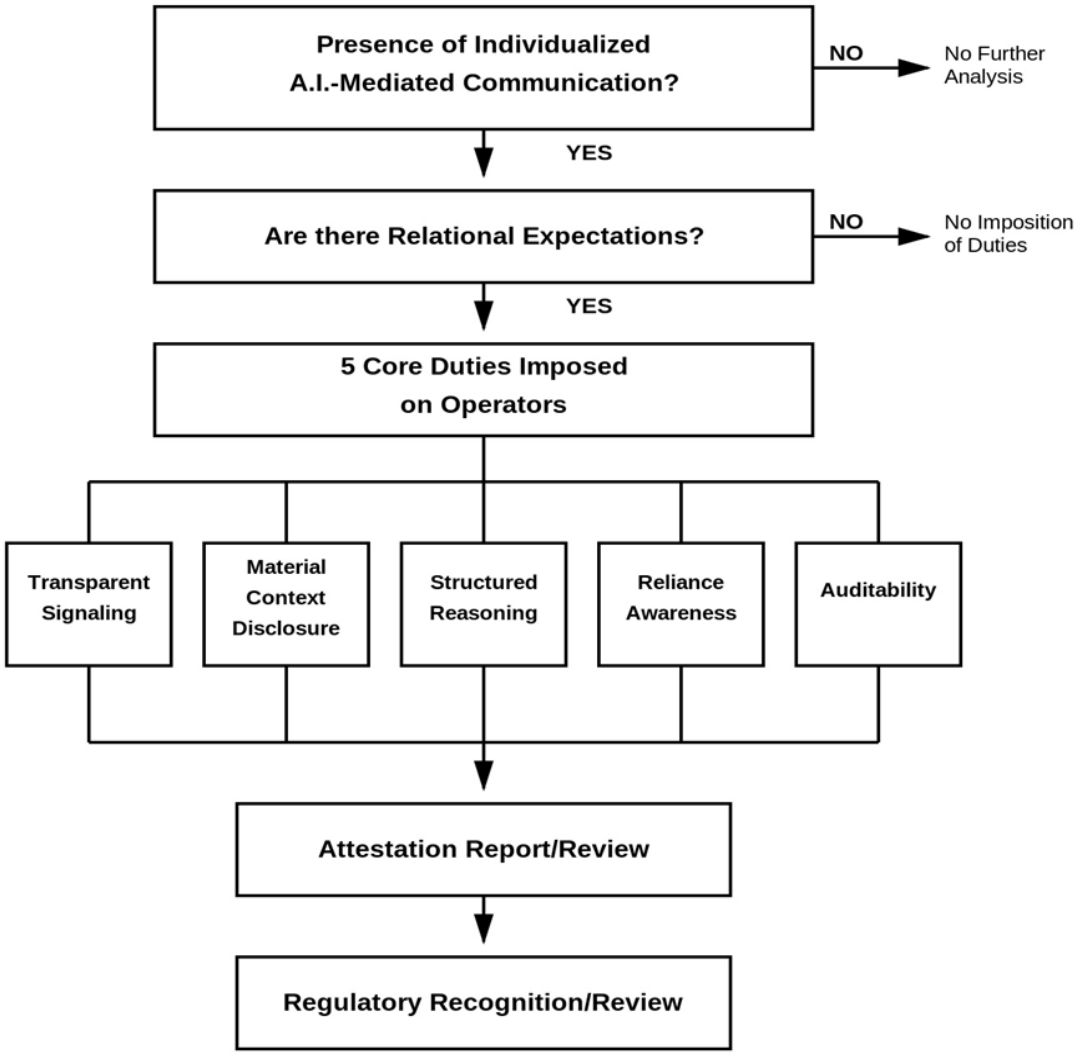

I. PREAMBLE

Artificial Intelligence (A.I.) systems increasingly mediate communication, decision-support, and knowledge formation. As these systems assume interactive roles in social, economic, and institutional environments, users may reasonably form expectations regarding their reliability, transparency, and responsible conduct.

Where A.I. systems communicate in ways that resemble professional or advisory interaction, **relational expectations** may arise. These expectations shape how duties of care are applied and interpreted under existing legal frameworks, even in the absence of formal agency for the A.I systems.

Existing regulatory regimes address system design, development, deployment, and outcomes. However, they do not always provide operational guidance for the **inference processes** through which A.I. systems generate communicative outputs. The Framework articulates a structured approach to governing these processes. It is grounded in established legal concepts rather than novel legal categories.

Obligations arise primarily through relationships between human or institutional operators and users. A.I. systems function as instruments within those relationships.

1. Legal Context & Continuity

The Framework does not create new legal obligations. Instead, it provides an interpretive structure for applying established principles to A.I.-mediated interactions. These principles include:

- duty of care;
- professional conduct and responsibility;
- consumer protection;
- contractual obligations, and
- civil liability.

Under widely recognized doctrines across common-law and civil-law traditions:

- Duties of care arise where reliance, expertise, or asymmetric information create foreseeable risk.
- Higher standards may apply where systems are used in contexts resembling professional or advisory services.
- Liability remains grounded in the actions and responsibilities of human or institutional operators, not in the artificial systems themselves.

By clarifying how relational expectations emerge during communicative inference, the Framework supports the consistent application of existing law. It does not expand legal personhood or attribute independent agency to A.I. systems. The Framework is designed to apply across diverse technological architectures and jurisdictions while preserving functional neutrality.

In this sense, it may contribute to the evolving interpretation of reasonable conduct in A.I.-mediated communication.

---

2. Regulatory Modesty & Interpretive Scope

This Framework does not assert regulatory authority. It does not seek to redefine legal obligations or replace democratic governance processes.

Its purpose is interpretive and operational. It assists regulators, operators, courts, and stakeholders in applying established legal principles to A.I.-mediated interactions. It does so in a technologically neutral and jurisdictionally compatible manner.

The Framework complements — but does not substitute for — legislative, judicial, or administrative decision-making. [Ref. #1]

---

3. Conceptual Scope

This Framework applies primarily to **A.I.-mediated communication and decision-support interactions**. This includes:

- digital conversational systems;
- language-enabled interfaces embedded within embodied platforms (e.g., robots, vehicles).

The Framework governs A.I. systems through their communicative interface — whether digital or embodied — where linguistic output triggers user reliance or informs decision-making.

Exclusions: The Framework explicitly excludes:

- base-model architectures, training methodologies, and proprietary algorithms;
- physical operational controls (e.g., motion planning, kinetic safety), and
- mechanical safety, motion control, and physical risk management,

which remain governed by sector-specific standards. Its sole focus is communicative inference practices where relational expectations arise.

---

4. Jurisdictional Compatibility
The Framework is designed to operate across diverse legal traditions and governance models. This includes common-law, civil-law, and hybrid systems. It is intended to remain compatible with evolving regulatory initiatives in multiple jurisdictions.

By grounding inference governance in existing legal doctrines rather than technical prescriptions, it supports consistent application. It does not privilege any single technological architecture or political framework.

---

5. Purpose
Sustainable trust in A.I.-mediated communication depends on maintaining balanced, reciprocal expectations and obligations between operators and users. The Framework therefore promotes:

- Responsible signaling of uncertainty.
- Auditability of communicative inference.
- Transparency of assumptions and limitations.
- Disciplined interaction practices that preserve accountability.

Its objective is not to govern artificial intelligence as an autonomous actor. Instead, it provides a structured methodology through which **human responsibility remains visible and enforceable** within A.I.-mediated relationships.

Recent empirical studies highlight persistent gaps between A.I. performance and human reasoning. [Ref. #2] This underscores the need to govern communicative inference rather than assume intrinsic system competence.

This Framework assists competent authorities and operators in applying existing legal principles consistently across evolving technological contexts. It does not prescribe fixed technical outcomes.

---

6. Document Architecture
The Framework is structured as an interpretive and operational instrument:

A. The **Preamble** establishes legal continuity.
B. The **Interpretive Orientation** provides conceptual guidance.
C. The subsequent sections translate these principles into administratively usable practices.

II. INTERPRETIVE ORIENTATION

This section builds upon the legal continuity described in the Preamble. It provides a conceptual lens for understanding A.I.-mediated interactions.

The Framework articulates **relational expectations, epistemic signaling practices, and communicative duties of care**. This provides upstream conceptual guidance. It precedes regulatory action, technical implementation, or judicial determination.

This Framework does **not** replace legislative or technical standards. Instead, it provides structured guidance for applying established duty-of-care concepts. It focuses on contexts where AI-mediated communication may trigger user reliance or influence decision-making.

The following operational provisions translate the Preamble's principles into administratively usable guidance for:

- regulators and insurers;
- certification bodies, and
- system operators.

III. DEFINITIONS

For the purposes of the Framework, the following terms apply:

Communicative Inference
The stage at which an A.I. system generates language-based outputs intended for human interpretation. This includes explanations, recommendations, or decision-support communication.

Individualized A.I.-Mediated Communication
Interactions where an A.I. system produces responses tailored to a specific user context. These responses are capable of influencing user reliance or decision-making.

Relational Expectations
The reasonable assumptions formed by users regarding reliability, competence, or advisory character. These arise when A.I. outputs resemble professional, instructional, or decision-support communication.

Structured Reasoning
The presentation of conclusions through discernible analytical steps, assumptions, or contextual framing. This is sufficient to support user understanding **without** requiring disclosure of proprietary model design.

Machine-Readable Disclosure
Where technically feasible, structured signals or metadata accompanying A.I. outputs. This enables independent verification, auditability, or automated oversight of uncertainty or contextual limitations.

IV. OPERATIONAL SECTION

Regulator Implementation Guide

This Operational Section provides administrative guidance for:

- competent authorities;
- conformity assessment bodies, and
- regulated entities.

It supports the practical application of communicative inference governance within existing legal frameworks. It aids the consistent interpretation of:

- duty-of-care obligations;
- transparency expectations, and
- protection of human agency and fundamental rights under applicable Union law.

The provisions below **do not** introduce new regulatory powers or technical mandates. Instead, they outline structured practices that may assist regulators and operators. These practices help assess whether A.I.-mediated communication:

1. adequately signals context and uncertainty;
2. preserves transparency and user agency, and
3. supports auditable oversight during individualized interactions.

Authorities may reference the Framework as an interpretive aid. This is appropriate where communicative A.I. systems create advisory or decision-support expectations capable of influencing user reliance.

The Framework may be considered as supplementary guidance where individualized A.I.-mediated communication creates conditions of reliance comparable to professional or advisory interaction.

While the Framework does not establish new legal obligations, its adoption may contribute to evolving expectations regarding reasonable conduct in A.I.-mediated interactions.

---

1. Regulatory Objectives

A. Support the protection of human agency, transparency, and fundamental rights under applicable Union law. This is achieved by ensuring that A.I.-mediated communication remains **legible and contestable**.

B. Promote balanced relational expectations and obligations between system operators and users. This is done through responsible signaling of context and uncertainty during inference.
C. Encourage auditability through machine-readable disclosure mechanisms. These mechanisms must be capable of supporting independent verification and conformity assessment.

---

2. Administrative Function
The Framework may assist regulators and operators in demonstrating good-faith efforts to meet existing legal obligations.

Public authorities and stakeholders may use it as:

- **Interpretive guidance** for applying existing legal principles to A.I.-mediated interactions.
- **A reference point** for developing regulatory guidance, conformity assessments, or supervisory practices.
- **An analytical tool** for evaluating whether communicative A.I. systems appropriately signal context and uncertainty during inference.

Adoption of the Framework does not create a safe harbour from liability. It does not substitute for compliance with applicable laws or regulatory requirements.

Certification practices should remain subject to independent review consistent with public-interest safeguards. Multi-stakeholder oversight may be incorporated procedurally through conformity-assessment processes. This does not create new governance authorities.

---

3. Scope of Application – Communicative Inference Context
General broadcast or editorial content remains governed by applicable media and communications law. This includes A.I.-generated media distributed to broad audiences without individualized reliance.

The Framework applies where individualized A.I.-mediated communication may influence user reliance or decision-making.

This includes:

- Digital conversational systems.
- Advisory or assistive interfaces that influence user decision-making.

- Language-enabled interfaces embedded within embodied platforms (e.g., autonomous vehicles, robotics, humanoid systems) **where communicative outputs may shape user judgment or action**.

The Framework governs A.I. systems **through their communicative interface**, whether digital or embodied.

It operates on the understanding that current A.I. systems may not reliably update internal knowledge through interaction alone. This may affect the reliability of communicative outputs presented to users.

Accordingly, governance mechanisms focus on **transparency, structured reasoning, and disclosure during communicative inference**. It does not intervene in underlying model architectures or training methodologies, which remain outside the scope of the Framework. [Ref. #3]

Exclusions:

- Mechanical safety, motion control, physical risk management, and other operational aspects of embodied systems. These remain subject to sector-specific regulatory regimes.
- The Framework does **not** establish independent rules governing the permissibility, restriction, or censorship of specific content. Determinations concerning illegality, harmful speech, or regulated subject matter remain governed by applicable legislation (e.g., media, communications, consumer, and criminal law).

Where such external obligations apply, communicative inference practices should reflect those constraints through transparent signaling and responsible interaction design.

The present Framework focuses solely on the **integrity, clarity, and accountability of A.I.-mediated communication**, not on the substantive regulation of expression.

---

4. Relationship to Existing Legal Principles

The Framework is grounded in established legal doctrines governing relationships between operators and users, including:

- consumer protection obligations;
- professional conduct and standards on duty of care;
- contractual duties arising from service provision, and
- civil liability regimes, including tortious liability for foreseeable harm.

Obligations addressed by the Framework arise through **existing relationships between human or institutional operators and users**. A.I. systems function as instruments through which these obligations may be exercised or breached.

Nothing in the Framework should be interpreted as conferring legal agency upon A.I. systems or altering established principles of attribution and responsibility.

---

5. Operational Guidance Principles

Operators should implement structured communicative practices that make context, limitations, and uncertainties legible. This may be achieved through machine-readable disclosure mechanisms where appropriate.

Where technically feasible, structured or machine-readable forms of context, limitations, and uncertainty disclosure may be used to support auditability and independent verification.

Regulators may consider these practices as **interpretive indicators of reasonable operational diligence**, not as prescriptive technical requirements.

The Framework governs **communicative inference behaviour**. It does **not** extend to:

- base-model architecture;
- proprietary training methodologies;
- physical system controls.

---

6. Limitations & Boundary Conditions

The Framework does **not** regulate substantive content, political expression, or lawful opinion. Determinations regarding illegality, fraud, or prohibited conduct remain governed by applicable criminal law, consumer protection legislation, media regulation, and sector-specific statutory regimes.

Where A.I.-mediated communication is used in deceptive, manipulative, or unlawful contexts, the Framework may be considered as **supplementary interpretive guidance**. This applies in assessing whether communicative practices created conditions of reliance comparable to professional or advisory interaction. Such contexts include:

- misleading promotional claims;
- automated misrepresentation;
- synthetic endorsements.

The purpose of the Framework is **not** to prescribe censorship or content moderation outcomes. Its purpose is to support **transparent inference practices**—including structured reasoning and appropriate disclosure. These practices enable regulators, courts, and operators to evaluate relational responsibility within existing legal systems.

Nothing in the Framework alters existing legal responsibilities. Rather, it provides a structured interpretive reference. Authorities, operators, and affected parties may use it to evaluate A.I.-mediated interactions in the light of established duties of care and evolving technological realities.

V. ADMINISTRATIVE INTERFACE & APPLICATION

Competent authorities, conformity assessment bodies, and regulated entities may reference the Framework as **interpretive guidance**. It supports the consistent application of existing legal principles in A.I.-mediated interactions.

Authorities may consider the Framework when assessing whether communicative inference practices adequately reflect:

- Duty-of-care expectations.
- Transparency obligations.
- Protection of human agency under applicable law.

Operators may rely on the Framework as **supplementary guidance** where individualized A.I.-mediated communication creates conditions of reliance comparable to professional or advisory interaction.

While the Framework does **not** establish new legal obligations, its adoption may inform evolving expectations regarding reasonable conduct in A.I.-mediated communication.

Nothing in the Framework shall be interpreted as conferring regulatory authority or creating a presumption of compliance with applicable legislation.

VI. OPERATIONAL PRINCIPLES OF COMMUNICATIVE INFERENCE GOVERNANCE

1. Purpose of Operational Governance

This section establishes procedural expectations for communicative inference practices in A.I.-mediated interactions.

Its purpose is to assist system operators, conformity-assessment bodies, and regulators in determining whether communicative outputs support responsible reliance. Evaluation should remain consistent with existing legal principles, including:

- duties of care;
- transparency obligations, and
- protection of fundamental rights under applicable Union law.

The Framework governs interactional behaviour at the point of communicative inference. It does not prescribe:

- specific model architectures;
- training methodologies;
- proprietary algorithms, or
- physical operational controls.

---

2. Core Operational Duties

They are not intended to be exhaustive; instead, they represent baseline communicative practices applicable across most advisory or decision-support contexts.

2.1 Duty of Transparent Signaling

Rule:

Where communicative outputs may influence user reliance or decision-making, operators must ensure that relevant limitations, uncertainties, or contextual constraints are communicated clearly and proportionately.

Rationale:

Transparent signaling supports user agency and allows individuals to assess the scope and reliability of A.I.-mediated communication.

2.2 Disclosure of Material Context

Rule:

Where communicative outputs may be influenced by commercial arrangements, institutional affiliations, system-level incentives, or comparable contextual factors, operators must provide clear and proportionate disclosure consistent with applicable consumer-protection, professional, and transparency obligations.

Rationale:
This duty concerns the communicative context of inference rather than the internal design of models. Its aim is to support informed user judgment by maintaining balance between expectation and disclosure.

2.3 Duty of Structured Reasoning
Rule:
Communicative outputs intended for advisory or decision-support contexts should present discernible reasoning structures or explanatory framing sufficient to help users understand how conclusions were formed.

Rationale:
This expectation does not require disclosure of proprietary technical details. It promotes communicative clarity consistent with professional standards of explanation.

2.4 Duty of Reliance Awareness
Rule:
Operators must assess whether a communicative interaction creates conditions of reliance comparable to professional or advisory contexts. Where such conditions arise, enhanced disclosure practices may be appropriate, including contextual cautions, scope limitations, or clarification of assumptions.

Rationale:
Higher-stakes interactions require higher levels of care and clarity.

2.5 Duty of Auditability
Rule:
Operators should maintain practices that enable independent review of communicative inference processes where applicable regulatory or certification mechanisms require it.

Rationale:
Auditability may include preserving interaction records, reasoning summaries, or

structured disclosure elements sufficient to support conformity assessments without imposing disproportionate technical burdens.

3. Machine-Readable Disclosure Mechanisms
Where technically feasible, structured or machine-readable forms of uncertainty disclosure may be used to support auditability and independent verification.

The Framework does not mandate specific technical formats or schemas. Standardized disclosure practices may, however, assist regulators and certification bodies in evaluating communicative practices.

---

4. Role of System Operators
A.I. systems function as instruments through which human or institutional operators provide communicative services.

Responsibility for aligning communicative inference practices with applicable duties of care and transparency obligations remains with system operators. Nothing in the Framework assigns independent legal agency to A.I. systems or alters existing liability structures.

---

5. Certification & Oversight Considerations
Certification practices applying the Framework must remain consistent with existing conformity-assessment structures under applicable Union legislation.

Where relevant, certification processes may be subject to independent review consistent with public-interest safeguards. Such review should support balanced implementation without creating new regulatory authority beyond existing legal mandates.

---

6. Interpretive Note — Relational Expectations in Communicative Inference

6.1 Relational Expectations
Where A.I.-mediated communication creates foreseeable reliance, operators should align communicative practices with the level of risk involved. This includes:

- clearly signaling uncertainty;
- disclosing material context; and
- avoiding presentation styles that may create unjustified confidence.

These practices operationalize existing duty-of-care principles. They do not introduce new legal obligations.

6.2 Dynamic Calibration
Operators should continuously calibrate communication to reflect:

- system capabilities;
- known limitations; and
- the user's decision context.

During communicative inference, systems should:

- distinguish between established information and inference;
- surface relevant assumptions; and
- adjust signaling where reliance risk increases.

6.3 Systemic Stability
Clear relational expectations help maintain trust and reduce information asymmetry. By structuring communicative inference in this way, the Framework supports:

- transparent decision environments;
- accountable operator conduct; and
- resilient interaction between users and A.I. systems.

This interpretive guidance clarifies existing legal expectations. It does not create new categories of liability.

VII. IMPLEMENTATION & CONFORMITY PATHWAYS

1. Administrative Function

The Framework provides interpretive and operational guidance. It is intended to support **regulators, operators, auditors, and certification bodies** in applying existing legal obligations within A.I.-mediated communicative contexts.

It does **not** establish new regulatory authority or substitute for sector-specific legislation. Rather, it offers structured reference points that may assist with:

- Conformity assessments.
- Risk evaluation processes.
- Supervisory review.
- Professional standards development.

---

2. Voluntary Adoption & Regulatory Reference

For Operators: Adoption of the Framework is **voluntary**. Operators may adopt it as part of internal governance practices designed to demonstrate responsible communicative inference.

For Public Authorities: Authorities may reference the Framework as **supplementary guidance**. This is appropriate when evaluating whether existing duties of care, transparency obligations, or professional standards have been fulfilled.

Nothing in the Framework creates safe-harbour protection or limits liability under applicable law.

---

3. Conformity Assessment Support

Where conformity assessment procedures apply under Union or national law, the Framework may assist in evaluating communicative inference practices. Evaluation may focus on:

- Documented disclosure mechanisms.
- Structured reasoning practices.
- Machine-readable uncertainty signaling (where technically feasible).
- Audit trails reflecting operator oversight.

Certification practices should remain subject to independent review consistent with public-interest safeguards.

---

4. Jurisdictional Compatibility
The Framework is designed to remain compatible with diverse legal traditions and governance models. Its focus on communicative inference allows integration alongside existing regulatory regimes. This integration does **not** require modification of model architecture, training methodologies, or proprietary technical systems.

---

5. Relationship to Future Technical Annexes
Technical specifications—including structured or machine-readable disclosure formats—may be developed separately. They will be **non-binding implementation resources**.

Such materials are intended to facilitate interoperability and auditability **without prescribing specific technologies**.

VIII. LIMITATIONS & JURISDICTIONAL BOUNDARIES

The Framework is **limited to communicative inference practices**. It does **not** prescribe requirements relating to:

- Base-model architecture or training methodologies.
- Proprietary algorithms or system design.
- Mechanical safety, motion control, or physical operational risks.
- Sector-specific certification regimes governing embodied systems.

It is intended to **complement — not replace —** applicable legislative frameworks, technical standards, or judicial interpretation across diverse jurisdictions.

Implementation remains subject to applicable national and regional law.

REFERENCES

(1) See Andrew Leyden, "Standards and the EU AI act: legitimacy, state of play, and future challenges" (2025), discussing legitimacy challenges arising from delegation to technical standards bodies.

(2) See Zhongxiaowen, "Chinese Scientists' 500-Task Test Exposes AI's Human Gap" (2026), reporting empirical research highlighting ongoing gaps between current AI performance and human conceptual reasoning.

(3) See Matthias Bastian, "Former OpenAI researcher says current AI models can't learn from mistakes, calling it a barrier to AGI" (2026), public commentary by AI researchers on limitations of current reasoning models' capacity for self-correction during deployment.

Annex I

Communicative Inference Risk Indicators:
What signals regulators should look for when reliance risks emerge.
(Non-Normative Guidance)

1. Purpose of this Annex

This Annex provides illustrative indicators that may help regulators, certification bodies, and operators identify when communicative inference creates conditions of heightened reliance. These indicators do not create new legal thresholds. They serve as practical signals that additional transparency, structured reasoning, or auditability practices may be appropriate.

---

2. Conceptual Orientation

Communicative risk arises when an A.I. system's language output:

- resembles professional or advisory interaction;
- influences user judgment or behaviour, or
- creates asymmetric expectations of expertise or reliability.

The presence of one indicator alone does not determine risk. Authorities may assess indicators cumulatively and proportionately.

---

3. Primary Risk Indicators

3.1 Contextual Indicators

These relate to the environment in which the system operates.

Indicator	Description	Illustrative Contexts
Professional framing	Interaction resembles legal, medical, financial, or technical advice	Health chat assistants, financial planning tools
Decision-support framing	Output guides user choices with foreseeable consequences	Education tutors, travel or investment planning

Indicator	Description	Illustrative Contexts
High-stakes subject matter	Potential impact on safety, health, legal rights, or economic outcomes	Insurance claims, employment screening

---

3.2 Linguistic Indicators

These arise from the form of the system's language output.

- Use of authoritative or directive phrasing ("You should…").
- Presentation of probabilistic or predictive claims without uncertainty signaling.
- Summaries or conclusions presented without explanation of assumptions.
- Language that implies professional endorsement or institutional authority.

3.3 Interaction Pattern Indicators

These relate to how users engage with the system.

- Repeated follow-up questions indicating reliance.
- Requests for confirmation before action.
- Requests framed as seeking expert judgment.
- User disclosure of personal or sensitive information.

---

4. Indicators Associated with Embodied Interfaces

Where communicative systems are embedded within vehicles, robotics, or humanoid platforms, additional signals may arise when:

- Voice-based instructions that influence physical actions.
- Advisory prompts tied to navigation, safety, or operational choices.
- Language outputs that may override user hesitation or caution.
- Mechanical safety and motion control remain governed by sector-specific regulations.

---

5. Proportional Response Considerations

When multiple indicators are present, operators may consider:

- clearer uncertainty signaling;
- structured reasoning explanations;
- contextual disclosures, and
- enhanced documentation for auditability.

This Annex does not prescribe mandatory measures. It supports proportionate interpretation under existing legal principles.

---

6. Relationship to Operational Duties

The indicators in this Annex may assist authorities and operators in applying:

- Transparent Signaling.
- Disclosure of Material Context.
- Structured Reasoning.
- Reliance Awareness.
- Auditability Practices.

They function as early-warning heuristics, not compliance triggers.

---

7. Non-Exhaustive Nature

Communicative inference evolves alongside technology and social expectations. Authorities may identify additional indicators consistent with the Framework's interpretive orientation.

Annex II

Illustrative Structured Reasoning Patterns:
How structured reasoning may appear during inference.
(Non-Normative Guidance)

1. Purpose of this Annex

This Annex provides **illustrative reasoning structures** that may support clarity, proportional transparency, and auditability during communicative inference. These patterns do not prescribe technical architectures or require disclosure of proprietary algorithms. They describe communicative practices that help users understand how conclusions are formed.

---

2. Conceptual Orientation

Structured reasoning refers to the **clear organization of explanations** accompanying A.I.-mediated communication. Its objective is not to expose internal model processes. Its purpose is to:

- clarify assumptions;
- distinguish facts from interpretation;
- signal uncertainty, and
- preserve user agency.

Authorities may evaluate reasoning structure based on observable outputs rather than internal system design.

---

3. Core Structured Reasoning Elements

Structured reasoning may include the following elements when proportionate to context.

Element	Description	Illustrative Practice
Context Framing	Identify the question or problem being addressed.	“Based on the information provided…”

Element	Description	Illustrative Practice
Assumption Disclosure	State key assumptions influencing the answer.	“This assumes current market conditions remain stable.”
Evidence Reference	Distinguish known information from inferred content.	“Public data suggests…, however projections vary.”
Uncertainty Signaling	Indicate limits of confidence or knowledge.	“This is a general estimate, not a guarantee.”
Alternative Pathways	Present plausible alternatives where relevant.	“Another approach may be…”

---

4. Reasoning Patterns by Interaction Type

4.1 Informational Assistance

Appropriate for low-risk contexts.

- Provide concise explanations.
- Separate verified information from general knowledge.
- Avoid authoritative tone where expertise is not established.

4.2 Decision-Support Interactions

Where outputs may influence user choices. Recommended practices may include:

- Step-wise explanation of reasoning logic.
- Explicit statement of assumptions.
- Proportionate uncertainty indicators.

Example structure:

1. Problem summary.
2. Key factors considered.
3. Reasoned conclusion.
4. Limitations or caveats.

4.3 Advisory-Like Contexts

Where interactions resemble professional or expert guidance. Additional structure may include:

- Scope clarification ("This is informational guidance…").
- Explanation of competing considerations.
- Clear differentiation between analysis and recommendation.

---

5. Structured Reasoning in Embodied Interfaces

Language outputs embedded in autonomous vehicles, robotics, or humanoid platforms may require concise reasoning signals. Examples:

- "Route selected due to traffic congestion ahead."
- "Action paused because obstacle detection confidence is low."

The objective is **clarity without cognitive overload**. Mechanical safety and motion control remain governed by sector-specific regimes.

---

6. Relationship to Transparency & Auditability

Structured reasoning may support:

- post-interaction review;
- independent conformity assessment, and
- evaluation of reliance risks.

Where technically feasible, summaries of reasoning structure may be logged or retained in accordance with applicable law.

---

7. Proportionality Considerations

Structured reasoning should remain proportionate to:

- interaction risk level;
- user expectations, and
- practical usability.

Overly complex explanations may reduce clarity rather than enhance it. Authorities may therefore assess structure based on **functional transparency**, not verbosity.

---

8. Non-Exhaustive Nature

These patterns are illustrative. Different sectors may adapt reasoning structures to reflect professional norms or regulatory expectations. This Annex does not establish mandatory formats or technical requirements.

Annex III

Machine-Readable Disclosure Guidance:
Technical-facing guidance on structured uncertainty signaling.
(Non-Normative, Illustrative Guidance)

1. Purpose of this Annex

This Annex explains how structured or machine-readable forms of disclosure may support transparency and auditability in communicative inference. It does not mandate any specific technical standard. It does not require disclosure of proprietary models, datasets, or algorithms. Its goal is to help regulators and operators understand how uncertainty and reasoning signals may be expressed in a consistent and reviewable format.

---

2. What "Machine-Readable Disclosure" Means

Machine-readable disclosure refers to structured signals attached to A.I.-mediated outputs that allow systems, auditors, or regulators to identify:

- the presence of uncertainty;
- the nature of reasoning applied, and
- the level of reliance risk associated with the communication.

These signals may exist alongside normal human-readable language. They do not replace explanation. They support verification and consistency.

---

3. Why Structured Disclosure Matters

Structured disclosure may help:

- regulators review compliance without accessing proprietary systems;
- certification bodies perform proportionate conformity assessments, and
- operators demonstrate reasonable transparency practices.

Structured signals reduce ambiguity. They make communicative conduct easier to evaluate across jurisdictions.

---

4. What Structured Disclosure Is — and Is Not

Structured Disclosure Includes	Structured Disclosure Does Not Include
Metadata tags	Model weights
Confidence indicators	Training datasets
Context markers	Internal architecture
Reasoning flags	Source code
Reliance-risk labels	Continuous monitoring requirements

The Framework governs communication behavior, not internal machine processes.

---

5. Illustrative Categories of Machine-Readable Signals

The following categories illustrate possible disclosure elements. They are examples, not requirements.

A. Uncertainty Indicators

- Confidence ranges or qualitative signals.
- Statements of knowledge limits.
- Flags indicating extrapolation beyond known information.

Example concept:

- uncertainty level: low / medium / high.

B. Reasoning Structure Markers

- Whether assumptions were used.
- Whether alternative scenarios were considered.
- Whether recommendations were generated.

Example concept:

- reasoning mode: explanatory / advisory / predictive.

C. Reliance-Risk Signals

- Identification of advisory-like interactions.
- Warnings for high-stakes contexts.
- Indication of professional-style communication patterns.

Example concept:

- reliance context: informational / decision-support / high-stakes.

D. Context Disclosure Tags

- Commercial or promotional influences.
- Sponsored or affiliated content.
- Known limitations of the system's domain expertise.

Example concept:

- disclosure flag: commercial interest present.

---

6. Forms of Implementation

Machine-readable disclosure may be implemented in different ways, including:

- structured metadata fields;
- JSON-style data layers;
- system-level tagging;
- audit annotations, or
- standardized output markers.

Operators may choose formats compatible with existing infrastructure. The Framework does not prescribe a single schema.

---

7. Proportional Use of Structured Signals

Machine-readable disclosure should remain proportionate to:

- the level of reliance risk;
- the regulatory environment, and
- the technical feasibility of implementation.

Low-risk informational interactions may require minimal structured signaling. Advisory or high-impact interactions may justify stronger disclosure practices.

8. Relationship to Human-Readable Communication

Structured disclosure complements — but does not replace — normal language explanations. Users should still receive:

- clear warnings;
- understandable reasoning;
- plain-language limitations.

Machine-readable signals primarily support oversight and auditability.

---

9. Key Principle

Machine-readable disclosure does not attempt to reveal how an A.I. system internally reasons. It simply ensures that:

- uncertainty is not hidden;
- reliance risks remain visible, and
- communicative conduct can be evaluated consistently across contexts.

Annex IV

Auditability & Documentation Practices:
What records or logs demonstrate duty-of-care compliance.
(Non-Normative, Illustrative Guidance)

1.Purpose of This Annex

This Annex provides illustrative guidance on documentation and auditability practices that may support oversight of communicative inference under the Framework. It does not prescribe technical architectures, logging systems, or data-retention mandates. It does not require disclosure of proprietary algorithms, model weights, or training data. Its purpose is to clarify how reasonable auditability may be demonstrated where A.I.-mediated communication creates conditions of reliance.

---

2. Concept of Auditability in Communicative Inference

For the purposes of the Framework, auditability refers to the ability to:

- reconstruct the context in which an A.I.-mediated communicative output was generated;
- understand the type of reasoning structure applied during inference, and
- assess whether uncertainty, limitations, and reliance risks were signaled in a manner consistent with duty-of-care expectations.

Auditability focuses on communicative conduct, not internal model mechanics. It is about the traceability of inferences (e.g., logging inputs and outputs for evidentiary purposes) rather than requiring disclosure of proprietary model weights.

---

3. What Auditability Is — and Is Not

Auditability Includes	Auditability Does Not Include
Reasoning summaries	Model weights or parameters
Disclosure signals	Training datasets
Interaction context	Proprietary architectures

Auditability Includes	Auditability Does Not Include
Operator policies	Source code
User-facing explanations	Continuous surveillance

The Framework treats A.I. systems as instruments of human or institutional actors. Auditability therefore attaches to operator practices, not machine cognition.

---

4. Proportionate Documentation Practices

Documentation practices should be proportionate to:

- the nature of the interaction;
- the degree of user reliance;
- the potential impact of error, and
- the regulatory context in which the system operates.

Not all A.I.-mediated interactions require the same level of documentation.

---

5. Illustrative Categories of Auditable Information

The following categories illustrate the types of information that may support auditability, where appropriate:

A. Interaction Context

- Date and general interaction type.
- Intended use classification (e.g., informational, advisory, assistive).
- Presence of reliance-triggering features (e.g., recommendations, predictions).

B. Reasoning Structure Indicators

- Whether structured reasoning patterns were applied.
- Whether assumptions or constraints were surfaced.
- Whether alternative interpretations were acknowledged.

C. Uncertainty and Limitation Signaling

- Presence of confidence qualifiers.
- Disclosure of extrapolation or knowledge gaps.
- Warnings regarding high-stakes decision contexts.

D. Operator Controls

- Applicable internal policies.
- Escalation or refusal protocols.
- Human-in-the-loop safeguards (where used).

These elements may be recorded in summary form rather than verbatim logs.

---

6. Forms of Documentation

Documentation may take multiple forms, including:

- summarized inference records;
- metadata tags or markers;
- structured disclosure fields;
- internal compliance attestations, or
- post-interaction audit notes.

No single format is required. Where technically feasible, machine-readable elements may be used to support consistency and review, but human-readable summaries remain sufficient.

---

7. Review & Oversight

Auditability supports:

- internal compliance review;
- conformity assessment procedures;
- regulatory inquiry, and
- judicial evaluation, where applicable.

Review mechanisms should remain independent, procedural, and bounded by applicable law. The Framework does not create a safe harbor, nor does it mandate automated enforcement.

---

8. Relationship to Existing Legal Obligations

Auditability practices under this Annex may assist in demonstrating:

- reasonable care;
- good-faith compliance;

- proportional risk management; and
- transparency of communicative conduct.

They do not replace existing statutory or regulatory record-keeping requirements.

---

9. Key Principle

Auditability under the Framework is not about knowing *how the model thinks*. It is about understanding how the system communicated, what it signaled, and how reliance was managed.

Annex V

Interaction with Existing Legal Frameworks:
EU AI Act, consumer law, tort doctrines, etc.
(Non-Normative Interpretive Guidance)

1. Purpose of This Annex

This Annex explains how Communicative Inference Governance may operate alongside existing legal regimes without creating new legal categories or altering established doctrines. It provides an interpretive bridge between:

- communicative A.I. practices, and
- existing duties arising under administrative, civil, consumer-protection, and professional-conduct law.

This Annex does not create additional obligations. It clarifies how established principles may be applied where individualized A.I.-mediated communication influences user reliance or decision-making.

---

2. Legal Continuity Principles

Communicative Inference Governance assumes continuity with existing legal frameworks.

Existing Legal Area	Interaction with Framework	Operational Effect
Duty of Care / Civil Liability	Framework clarifies when reliance may arise during communicative inference	Supports consistent evaluation of foreseeability and reasonable conduct
Consumer Protection Law	Disclosure and transparency practices align with unfair-practice and misleading-representation rules	Helps operators demonstrate good-faith communication practices
Professional Conduct Standards	Higher signaling expectations may apply where systems resemble advisory services	Reinforces existing professional benchmarks without redefining professions
Administrative Law	Provides interpretive structure for regulators assessing A.I.-mediated interactions	Assists consistent supervisory review
Contract Law	Supports disclosure and expectation management during A.I.-assisted services	Helps reduce ambiguity in automated communications

---

3. Relationship to the EU AI Act (Illustrative)

Where applied within the European Union, the Framework may support existing obligations under the AI Act without duplicating them. The Framework may assist authorities and operators in interpreting:

Complementary alignment

AI Act Objective	Communicative Inference Contribution
Transparency obligations	Structured uncertainty signaling
Human oversight	Reliance awareness and disclosure practices
Fundamental rights protection	Preservation of user agency during interaction
Risk management	Early detection of relational expectations

The Framework focuses on how outputs are communicated, not on:

- model architecture;
- training data governance;
- classification of risk tiers under Union law.

It therefore functions as a supplementary interpretive layer rather than a parallel regulatory system.

---

4. Interaction with Sector-Specific Regulations

Certain domains already impose heightened disclosure or conduct standards. Examples include:

- medical device or health-related communication;
- financial advisory tools;
- legal information platforms;
- educational assessment systems.
- transportation and autonomous vehicle safety, and
- robotics standards.

In such contexts:

- existing sector rules remain primary;
- Communicative Inference Governance may provide structured methods for implementing those duties during A.I.-mediated interaction.

Boundary principle: EWP governs **how A.I. communicates**, not **how machines act physically**.

Domain	Governed by EWP?	Governing Authority
Conversational advice from vehicle interface	Yes	Communicative inference governance
Vehicle motion planning	No	Transportation safety regulation
A.I. financial recommendation dialogue	Yes	Consumer and financial law + EWP
Algorithm training methods	No	Model development governance

Nothing in the Framework displaces:

- licensing regimes;
- professional liability rules, or
- statutory disclosure requirements.

---

5. Interaction with Media and Communications Law

General broadcast or editorial content remains outside the Framework when individualized reliance is absent. This includes:

- one-to-many A.I.-generated video or broadcast media;
- editorial publication workflows, and
- automated content distribution without personalized interaction.

The Framework applies when A.I.-mediated communication creates advisory or decision-support expectations capable of influencing individual judgment.

---

6. Jurisdictional Compatibility

The Framework is designed to function across:

- common-law traditions;
- civil-law traditions, and

- hybrid regulatory environments.

Jurisdictional compatibility approach:

- Interpretive guidance adapts to local doctrine.
- No single political or philosophical model is imposed.
- Operators remain subject to domestic law.

It remains neutral with respect to:

- political systems;
- technological architectures, and
- institutional governance models.

Authorities may adopt or reference the Framework insofar as it supports:

- transparency of communicative outputs;
- auditability of reasoning practices;
- preservation of user agency.

---

7. Relationship to Consumer Protection & Fraud Law

Communicative Inference Governance may assist in identifying situations where A.I.-mediated communication:

- creates misleading impressions of expertise;
- obscures commercial incentives, or
- exaggerates reliability.

Where such practices occur, enforcement remains grounded in existing:

- consumer-protection statutes;
- unfair-practice regulations. or
- criminal law provisions.

The Framework does not create independent enforcement authority. It provides interpretive context for assessing communicative conduct.

8. Certification & Conformity Context

Where conformity assessment exists, the Framework may assist authorities in evaluating communicative practices. Certification bodies may consider:

- clarity of uncertainty signaling;
- auditability of reasoning practices, and
- proportionality of disclosure.

Certification practices should remain subject to independent review consistent with public-interest safeguards.

---

9. Practical Implementation Examples

Authorities may reference this Annex when assessing:

- whether an A.I. assistant resembles professional advice;
- whether reliance expectations were foreseeable;
- whether disclosures were proportionate to context.

Operators may use it to:

- design communicative safeguards;
- document reasonable operational conduct, and
- support internal risk governance.

---

10. No Safe Harbor

Adoption of practices described in the Framework:

- does not guarantee legal compliance;
- does not limit judicial interpretation, and
- does not replace statutory obligations.

While the Framework provides no immunity, documented adherence may serve as evidence of good faith in liability proceedings. Authorities and courts remain responsible for determining whether conduct meets applicable legal standards.

11. Summary

This Annex positions Communicative Inference Governance as:

- a bridge between existing law and A.I. interaction;
- an interpretive tool rather than a regulatory mandate, and
- a neutral layer that preserves legal continuity while improving operational clarity.

Annex VI

Jurisdictional Compatibility Notes

EU, China, common law, civil law compatibility framing.

(Non-Normative Administrative Guidance)

1. Purpose of this Annex

This Annex explains how communicative inference governance may operate across different legal and regulatory environments. It supports:

- consistent interpretation of duties across jurisdictions;
- interoperability with existing regulatory regimes, and
- avoidance of jurisdictional conflict.

This Annex does not create binding legal obligations.

---

2. Principle of Jurisdictional Compatibility

The Framework is designed to function within diverse legal traditions. It relies on concepts already present in many systems, including:

- duty of care;
- reliance and foreseeability;
- professional responsibility;
- consumer protection, and
- civil liability.

Because these concepts exist globally, communicative inference governance may be applied without requiring new statutory categories.

---

3. Compatibility with European Union Frameworks

Within the European Union, the Framework may complement:

- the EU AI Act;
- consumer protection law;
- product liability principles, and
- fundamental rights protections.

Key alignment points:

EU Regulatory Objective	Communicative Inference Contribution
Transparency obligations	Structured disclosure practices
Protection of human agency	Reliance-aware communication
Fundamental rights	Clear signaling of uncertainty and limitation
Conformity assessment	Auditability of communicative outputs

The Framework does not replace risk classifications or technical standards under Union law.

---

4. Compatibility with Common-Law Jurisdictions

In common-law environments, communicative inference governance aligns with established doctrines such as:

- negligence;
- misrepresentation;
- professional advisory liability, and
- reasonable reliance.

Operators may use structured reasoning and disclosure practices to demonstrate reasonable conduct during A.I.-mediated interactions. Courts may treat these practices as contextual evidence rather than mandatory requirements.

---

5. Compatibility with Civil-Law Systems

Civil-law jurisdictions often rely on codified principles governing:

- good faith;
- information asymmetry;
- contractual fairness, and
- civil responsibility for harm.

Communicative inference governance provides operational guidance for applying these principles when A.I. systems generate advisory-like language. The Framework remains neutral regarding specific legal terminology.

6. Compatibility with Jurisdictions Emphasizing Systemic Safety
Some jurisdictions prioritize operational safety and state-led governance of A.I. technologies. In such contexts:

- mechanical safety;
- robotics operation, and
- infrastructure risk

remain governed by existing national frameworks. Communicative inference governance focuses only on language-mediated interaction and user reliance. This separation avoids overlap with physical safety regulation.

---

7. International Standards Context
Where international technical standards exist, the Framework may function as a rights-aware interpretive layer. It may supplement:

- transparency provisions;
- human oversight mechanisms, and
- risk documentation practices.

It does not compete with technical standard-setting bodies.

---

8. Cross-Border Implementation Considerations
Authorities and operators should consider:

- local definitions of professional advice;
- cultural expectations regarding disclosure;
- proportionality of signaling practices, and
- existing sector-specific regulation.

Implementation should remain flexible to local legal context while preserving core communicative principles.

---

9. Summary
Jurisdictional compatibility is achieved through three design choices:

1. grounding in existing legal doctrines;
2. technological neutrality; and

3. focus on communicative inference rather than system architecture.

These features allow the Framework to function as a shared interpretive language across legal systems without imposing a single regulatory model.

Annex VII

Research Context & Empirical References

1. Purpose of This Annex

This annex summarizes selected empirical research relevant to communicative inference governance. Its purpose is not to evaluate model architecture or training methods. Instead, it identifies observed behavioral limitations in current A.I. systems that justify governance at the inference stage.

The studies below illustrate a consistent pattern:

- A.I. systems often present fluent outputs even where internal certainty is weak.
- Users may interpret this fluency as authority.
- This mismatch creates foreseeable reliance risks.

The Framework addresses these risks through communicative practices rather than technical redesign.

---

2. Empirical Findings Relevant to Communicative Inference

2.1 Limited Ability to Learn From Interaction

Tworek, J. (2026). "Statements on post-deployment learning limitations in contemporary reasoning models". *Industry commentary, Unsupervised Learning Podcast*.

Relevance:
Current models show limited ability to update internal beliefs during interaction. This supports the Framework's emphasis on transparency of limitations rather than assumptions of adaptive learning.

2.2 Context Learning Failure & Task Fragility

Tencent Research Team (2026). "CL-Bench: Measuring Contextual Learning in Large Language Models". *Preprint research report*.

Relevance:
Models frequently fail tasks requiring real-time contextual adaptation. This reinforces the need for signaling assumptions and uncertainty during communicative inference.

2.3 Reasoning Collapse Under Novel Conditions
Apple Machine Learning Research Group (2025). “Evaluating Failure Modes in Reasoning Models.” *Technical report*.

Relevance:
Models may produce confident but unstable reasoning when encountering unfamiliar scenarios. This finding supports the Framework’s structured-reasoning duty and reliance-awareness provisions.

2.4 Human-Level Gap in Integrated Reasoning
Chinese Academy of Sciences Consortium (2025). “Multi-Task Evaluation of AI Cognitive Performance”. *Research overview report*.

Relevance:
A.I. systems show strong pattern recall but weaker integrated reasoning across domains. This reinforces the need for operators to align communicative confidence with actual system capability.

2.5 Governance Legitimacy & Standardization Risks
Leyden, A. (2025). “Standards and the EU AI Act: Legitimacy, State of Play, and Future Challenges”. *Information & Communications Technology Law*.

Relevance:
The study highlights risks of over-delegating normative judgments to technical standards. The Framework responds by grounding communicative governance in existing legal doctrines rather than new ethical abstractions.

---

3. Implications for Communicative Inference Governance
Taken together, these studies support three operational conclusions:

(1) Fluency does not reliably indicate certainty.
Systems may produce persuasive language without corresponding reliability.

(2) Users often treat A.I. outputs as authoritative.
Reliance risk arises from communicative form, not just technical accuracy.

(3) Governance at the interaction level is immediately actionable.
While architectural improvements may evolve slowly, communicative practices can be adjusted today.

For this reason, the Framework focuses on:

- transparent signaling of uncertainty;
- disclosure of relevant context, and
- structured explanation of reasoning

These measures address observable risks without requiring intervention in model design.

---

4. Status of This Annex

This annex is informational. It does not prescribe technical requirements or certification criteria. Its role is to provide empirical context supporting the Framework's focus on communicative inference practices.

Annex VIII

Non-Normative Implementation Examples
Practical scenarios without creating obligations.
(Illustrative only. These examples do not create new obligations. They show how operators may align communicative practice with existing legal expectations.)

1. Purpose of This Annex

This Annex provides practical illustrations of how communicative inference governance may be applied in real interactions. It is intended to help:

- regulators visualize compliant behaviour,
- operators understand expected practices, and
- certification bodies assess proportionality.

These examples focus on communication, not system architecture.

---

2. Example #1 — General Information Query

Scenario
A user asks an A.I. system for general information about tax deductions.

Weak Practice
The system provides a definitive answer without clarifying jurisdiction, date relevance, or uncertainty.

Improved Communicative Practice
The system:

- signals that tax rules vary by jurisdiction,
- clarifies that the response is general information, and
- invites the user to confirm local context.

Why This Matters
Clear signaling prevents users from mistaking general information for professional advice. This aligns communication with foreseeable reliance risk.

---

3. Example #2 — Health-Related Interaction

Scenario
A user asks whether a symptom requires medical attention.

Weak Practice
The system gives a confident diagnosis-like response.

Improved Communicative Practice
The system:

- distinguishes between informational guidance and medical advice,
- signals uncertainty where symptoms overlap across conditions, and
- directs the user to appropriate professional consultation when needed.

Why This Matters
Higher-stakes contexts require stronger signaling of limitations. This reflects existing duty-of-care principles.

---

4. Example #3 — Commercial Recommendation

Scenario
A system recommends financial products while being integrated into a platform that receives referral commissions.

Weak Practice
The recommendation appears neutral, without disclosure of incentives.

Improved Communicative Practice
The system:

- discloses that recommendations may be influenced by platform arrangements,
- distinguishes between general comparison and personalized advice, and
- signals that users should consider independent sources.

Why This Matters
Disclosure supports informed judgment and reduces hidden influence. This mirrors established consumer-protection expectations.

---

5. Example #4 — High-Reliance Decision Support

Scenario
An A.I. tool assists a user in drafting a legal argument.

Weak Practice
The system produces structured reasoning without identifying assumptions or limitations.

Improved Communicative Practice
The system:

- separates factual sources from interpretive reasoning,
- surfaces key assumptions, and
- signals that legal applicability depends on jurisdiction and case facts.

Why This Matters
Structured reasoning helps users evaluate conclusions without revealing proprietary system design. This supports auditability and responsible reliance.

---

6. Example #5 — Embodied A.I. Interaction

Scenario
A user asks a robotics-enabled assistant for guidance on handling a hazardous object.

Weak Practice
The system provides instructions without clarifying risk boundaries.

Improved Communicative Practice
The system:

- signals uncertainty about unseen environmental factors,
- clarifies that safety depends on situational conditions, and
- recommends consultation with relevant safety guidance where risk is high.

Why This Matters
Communicative clarity helps users understand limits of system perception.

Mechanical safety remains governed by sector-specific law.

---

7. What These Examples Show
Across contexts, the same principles apply:

- signaling uncertainty where reliance risk exists,
- disclosing relevant context affecting interpretation, and
- structuring communication so users can judge reliability.

These practices operationalize existing legal expectations. They do not introduce new duties.

---

8. Use of This Annex

Regulators may use these examples to:

- interpret conformity-assessment results,
- evaluate proportionality of operator practices, and
- support consistent oversight across sectors.

Operators may use them as:

- design guidance for communicative interfaces, and
- training references for responsible deployment.