



CORNER ITALIANO | OFFICIAL RESOURCES

Complete Glossary

Technical terms, operational concepts, and original proposals from the book

REFERENCE GUIDE

Aligned with the revised English interior · Documentary snapshot July 2026 · Resource version EN-US-R2.3

<https://www.corneritaliano.com/en/2050/resources/>

SCOPE A guide to the technical terms used in the book, explained through what they change in everyday life. The definitions clarify how the words are used in these pages. Standards, regulations, and specialist manuals remain the sources for technical or professional decisions.

Artificial intelligence (AI)

A set of techniques that allow a system to perform tasks associated with human capabilities, such as recognizing, predicting, translating, recommending, or generating. A convincing performance does not by itself demonstrate consciousness or establish that the system has experience or understanding.

Machine learning

An approach in which a model changes its parameters from examples instead of receiving only rules written one by one. What it learns depends on how the problem is formulated and on the data. The criteria used to evaluate it also shape the result.

Deep learning and neural networks

A neural network is a mathematical structure of connected units and parameters that can change during training. Deep learning uses networks with many layers to build useful representations of images, speech, language, and other signals. The name evokes the brain, but it does not mean that a small brain is hidden inside the computer.

Transformer

An architecture that became influential with the 2017 paper ‘Attention Is All You Need.’ It uses self-attention mechanisms to weigh relationships among positions in a sequence and allows many positions to be processed in parallel during training. It does not ‘read’ a text as a person does.

Large language model (LLM)

A neural model trained on large collections of text to estimate which tokens may follow the available context. During generation, it does not necessarily choose the most likely alternative every time. Instructions, context, selection method, and later training stages all matter. Fluency does not guarantee truth or understanding.

Token

A conventional unit into which a model divides text. A token may be a word, part of a word, punctuation, or another segment. In this book, the term is used only in its linguistic sense.

Training data and dataset

A dataset is a collection organized for analysis, training, or evaluation. Training data are the examples used to change a model’s parameters. Coverage, provenance, labels, and omissions reflect choices about what to observe and which examples to include or exclude. These choices influence both capabilities and errors.

Hallucination

A plausible but false or unsupported response produced with the same fluency as correct information. The term describes a type of output and does not imply that the system is having a perceptual experience.

Algorithm

A procedure made of steps or rules for transforming data and producing a result. Even when formalized, it incorporates choices about objectives, categories, thresholds, and exceptions.

Algorithmic bias

A systematic distortion that may arise from data, labels, objectives, or the context of use and may produce unequal effects among people or groups.

Profiling

The processing of data to evaluate or predict aspects of a person, such as preferences, behavior, reliability, or risk. Its lawfulness and limits depend on the context and the applicable rules.

Predictive surveillance

The use of data and models to anticipate behavior or risk and guide controls or interventions. It can turn a prediction into present treatment of a person. Its use requires defined limits and verification, and those affected must be able to contest it.

Unmeasured interval

A temporary space in which a person can hesitate or contradict themselves, and change without every variation immediately producing an operational profile. This space is compatible with the presence of technology, though it offers no guarantee of freedom. Continuous collection requires a defined harm and a proportionate purpose, together with a verifiable expiration date.

Distributed ledger or blockchain

A distributed ledger preserves shared state among several parties. A blockchain is one possible structure for doing so. It can help independent organizations coordinate without a single operator, but it does not make entered data true or eliminate trust. Digital identity, verifiable credentials, and selective disclosure do not require a blockchain.

Verifiable credential

An attestation composed of relevant attributes and issued by an accountable issuer. It is secured so that a verifier can check its provenance and integrity. Technical verification does not make a false fact true. The quality of the issuer, procedures, credential status, and the ability to correct or revoke it all matter.

Selective disclosure

The presentation of only the necessary attributes contained in a credential, while keeping the others hidden. It reduces the excess disclosure of information, but it does not by itself guarantee anonymity or prevent different interactions from being linked.

Zero-knowledge proof

A cryptographic proof that makes it possible to demonstrate possession of information or satisfaction of a condition without revealing the underlying data. It is useful in some contexts, but it does not replace sound privacy and accountability architecture.

Digital wallet

A tool with which a person can store digital credentials or attributes and manage how they are presented. It must address security, accessibility, recovery, revocation, delegation, and the availability of assistance. Personal control must remain compatible with access to assistance.

Digital identity

The set of means and attributes through which a person or organization is recognized in systems. It is not the whole person and is not identical to a single account. Identification, authentication, and authorization answer different questions.

Authentication and authorization

Authentication verifies factors associated with a claimed identity, such as a credential, a device, or several combined factors. Authorization establishes what that person or system may do in context. Passing the first does not automatically confer the second.

Interoperability and portability

Interoperability lets different systems exchange data and recognize functions or attestations under shared rules. Portability keeps data, identity, permissions, configurations, and logs usable outside the primary environment. It is incomplete without recovery, revocation, assistance, and bounded delegation. Neither one alone guarantees continuity.

Data minimization

The GDPR principle that personal data must be adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed.

Proof minimization

A design proposal in this book to demonstrate only what a decision requires, without handing over the person's entire biography. It extends the logic of proportionality to the design of digital verification. See also: Verifiable credential; Selective disclosure; The 2050 Test.

Cloud

A set of remote computing resources used to host data and services. The word suggests lightness, but the cloud depends on buildings, networks, contracts, jurisdictions, skills, and energy.

Data center

A physical facility that houses servers, networks, cooling systems, and electrical continuity. It gives digital infrastructure a material presence in a territory, supported by organizational arrangements.

API

An application programming interface that defines how two software systems can request data or actions from each other. APIs enable integration and, at the same time, create technical and economic dependencies.

Internet of Things (IoT)

A network of physical objects equipped with sensors, software, and connectivity that can collect data or receive instructions: medical devices, meters, vehicles, traffic lights, appliances, and many others.

Deepfake

A synthetic image, audio recording, or video that convincingly imitates a real person. Provenance and transformations can help verification, but no single technical signal establishes the whole truth by itself.

Epistemic fatigue

A condition in which the cost of reconstructing a conclusion grows to the point that a person suspends a judgment that remains open to revision or withdraws from participation. It differs from credulity or disinterest. Interfaces, access to sources, time, and the ability of institutions to make verification practicable also shape it.

Digital twin

A computational representation of an object, process, or system, updated with data so that possible states can be observed or simulated. It does not necessarily imply an agent capable of acting on behalf of a person.

Operational twin

A software agent proposed in this book, tied to a principal and to an observable Mandate. It operates in Assists, Proposes, Decides, or Executes mode only within declared tasks, data, duration, thresholds, and authority to stop. As a tutor twin, it exposes the steps of the work and any errors, along with possible alternatives, so that competence can grow and the need for assistance can diminish. Capability and reliability remain distinct from each other and from authority.

The 2050 Test

A non-compensatory decision method proposed in this book. It organizes, in this order, Purpose / Mandate / Evidence / Distribution / Harm / Redress / Exit. Each answer receives an answer status, Sufficient / Unresolved / Blocking, and leads to an organizational decision, Proceed / Limit and test / Do not delegate / Stop. It neither assigns a score nor certifies a product. A Blocking condition is not offset by the other benefits.

Operational mandate

A written scope that connects a principal to a mode, Assists / Proposes / Decides / Executes, and specifies tasks, data, people affected, locations, duration, volume, logs, thresholds, prohibited actions, and authority to stop. If the model, data, provider, scale, task, or responsible parties change, the Mandate must be narrowed or reviewed. Technical capability alone does not define this scope.

Technological sovereignty, strategic dependency, and infrastructural voice

Sovereignty is the relative ability to understand, choose, negotiate, and replace dependencies. It does not require autarky. A dependency becomes strategic when it reduces a capability before another practical path exists. Infrastructural voice gives those affected the concrete power to name harm, suspend, correct, export, and replace. Consultation without consequence does not provide that power.

Right to erasure, also called the right to be forgotten

Under the GDPR, the right to obtain erasure of personal data when the stated conditions apply, subject to limits and exceptions. The GDPR does not apply to data concerning deceased persons, and national rules differ. The post-mortem digital directive - express consent, scope, revocation during life, and shutdown - is a proposal in this book, not a uniform right already in force.

Authentic archive, memorial, post-mortem simulation, and impersonation

The archive preserves original traces. The memorial organizes them without speaking as the deceased person. The generative simulation produces new responses or representations from the data, while impersonation deceptively presents them as coming from the real person. Ownership of the files does not amount to a Mandate to generate a first-person voice. Neither does kinship or payment.

Verified comprehensibility

The ability of a person or independent body to reconstruct why a system produced an outcome and test it on both an ordinary case and an edge case. The verification must be usable before the consequence occurs. A generic explanation or a formally complete case file does not establish this ability. If the verification cannot change or stop the outcome, the Mandate must be narrowed.

Structural opacity

A form of opacity that emerges when many components, updates, agents, and organizations combine in such a way that no actor possesses the complete picture in time to control the outcome. Opening the code may help without resolving it. See also: Verified comprehensibility; Time to harm and control window; Oversight budget.

Capability, reliability, and authority

Three different questions: what a system can do, how consistently it does so under real conditions, and what it is permitted to do on someone's behalf. Capability and reliability do not by themselves create authority or require adoption or complete delegation.

Human oversight

Oversight is meaningful only when an identifiable person has readable evidence, time, competence, and authority to modify or stop the action. The handoff must remain on record, with the reasons and subsequent review. The formal presence or signature of an operator is not enough.

Meaningful recourse

The concrete ability to receive an understandable notice and identify and correct the relevant elements. A person must be able to reach, without obstacles, an office capable of suspending the effect and reopening the case. That office must communicate a reasoned outcome within a time frame compatible with the harm.

Residual harm and operational case file

An honest application of the Test does not promise zero harm. The case file retains value when it distinguishes original sources from human annotations and generated summaries. It must acknowledge what cannot be reconstructed and connect the evidence to remedial actions and restrictions. Its purpose is to support learning and correction. A case file cannot establish that control is perfect.

Time to harm and control window

Time to harm separates the error from its consequence. The control window is the time actually available to recognize an error, reconstruct what happened, and intervene. Oversight is meaningful only if the necessary operations fit within that window. Otherwise, it becomes nominal.

Oversight budget

The relationship between the volume and speed of the decisions requiring oversight and the human resources available to review them attentively. When the load persistently exceeds the capacity for verification, human intervention becomes a belated signature. See also: Human oversight; Verified comprehensibility; Degraded mode.

Continuity, degraded mode, and residual competence

Continuity preserves an essential function while also recording the costs transferred to users and staff. Degraded mode provides a deliberately narrower service, with declared limits, during an outage. Residual competence is the local knowledge and authority needed to activate it and manage exceptions. It also enables staff to reconcile the work when the system returns.

Exit test

A practical verification, before adoption and then periodically, that data, logs, configurations, and permissions can be exported, reconstructed, used in degraded mode, and restored to operation without the primary provider. The organization must be able to resume work while recording losses, then return without hidden dependence. A portability clause alone does not establish that ability.

Content provenance

Verifiable information about a piece of content's acquisition, origin, transformations, and approvals. These records distribute responsibility and help reconstruct a technical history. They do not by themselves certify sincerity or establish the context or truth of the content. Incomplete metadata do not automatically make a document false.

Living license

A proposal in this book for the use of a voice, image, or work in generative systems: consent bounded by purpose and duration, traceability of uses, compensation where applicable, the possibility of revocation, and treatment of derivative works. It does not describe a uniform license already recognized under current law. See also: Content provenance; Operational mandate; The 2050 Test.

Life cycle assessment (LCA)

A method for assessing environmental impacts across declared stages and boundaries, from raw materials to end of life. Results depend on the goal, functional unit, data, and scope and cannot be compared when these elements change without being disclosed.

Material metabolism

A way of reading digital infrastructure through the materials and work it requires, following their circulation and what remains. It covers energy, water, minerals, equipment, labor, emissions, and waste. The book proposes this lens without reducing it to a single metric. Comparisons require declared boundaries and units, together with the locations to which they apply. See also: Life cycle assessment; Water withdrawal and consumption; Technological sovereignty.

Water withdrawal and consumption

Withdrawal is water removed from a source. Consumption is the portion not immediately returned to the local environment. Confusing them makes it difficult to compare sites or technologies and to interpret seasonal differences.

Solely automated decision

In the sense of Article 22 of the GDPR, a decision made without meaningful human involvement that produces legal effects or similarly significantly affects a person. The applicable law defines its scope and sets the exceptions and safeguards.

High-risk AI system

A legal category under the AI Act that applies to systems identified according to the functions, contexts, and conditions specified by the regulation. It is not a generic synonym for a dangerous system and carries specific obligations that apply progressively.

General-purpose AI model (GPAI)

A model capable of performing a wide range of tasks and being integrated into different systems. Under the AI Act, it is a regulatory category distinct from the individual system that incorporates it.

Note on use

The definitions describe how the terms are used in the book and distinguish, where necessary, technical concepts, legal categories, narrative scenarios, and the author's proposals. Professional or regulatory decisions require consultation of the primary source and the applicable law.