



BSI Standards Publication

Information technology — Artificial intelligence (AI) — Transparency taxonomy of AI systems

National foreword

This British Standard is the UK implementation of EN ISO/IEC 12792:2025. It is identical to ISO/IEC 12792:2025.

The UK participation in its preparation was entrusted to Technical Committee ART/1, Artificial Intelligence.

A list of organizations represented on this committee can be obtained on request to its committee manager.

Contractual and legal considerations

This publication has been prepared in good faith, however no representation, warranty, assurance or undertaking (express or implied) is or will be made, and no responsibility or liability is or will be accepted by BSI in relation to the adequacy, accuracy, completeness or reasonableness of this publication. All and any such responsibility and liability is expressly disclaimed to the full extent permitted by the law.

This publication is provided as is, and is to be used at the recipient's own risk.

The recipient is advised to consider seeking professional guidance with respect to its use of this publication.

This publication is not intended to constitute a contract. Users are responsible for its correct application.

© The British Standards Institution 2025
Published by BSI Standards Limited 2025

ISBN 978 0 539 21369 0

ICS 01.040.35; 35.020

Compliance with a British Standard cannot confer immunity from legal obligations.

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 30 November 2025.

Amendments/corrigenda issued since publication

Date	Text affected

English version

Information technology - Artificial intelligence (AI) - Transparency taxonomy of AI systems (ISO/IEC 12792:2025)

Technologies de l'information - Intelligence artificielle
(IA) - Taxonomie pour la transparence des systèmes
d'IA (ISO/IEC 12792:2025)

Informationstechnologie - Künstliche Intelligenz -
Transparenz-Taxonomie von KI-Systemen (ISO/IEC
12792:2025)

This European Standard was approved by CEN on 29 August 2025.

CEN and CENELEC members are bound to comply with the CEN/CENELEC Internal Regulations which stipulate the conditions for giving this European Standard the status of a national standard without any alteration. Up-to-date lists and bibliographical references concerning such national standards may be obtained on application to the CEN-CENELEC Management Centre or to any CEN and CENELEC member.

This European Standard exists in three official versions (English, French, German). A version in any other language made by translation under the responsibility of a CEN and CENELEC member into its own language and notified to the CEN-CENELEC Management Centre has the same status as the official versions.

CEN and CENELEC members are the national standards bodies and national electrotechnical committees of Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Norway, Poland, Portugal, Republic of North Macedonia, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Türkiye and United Kingdom.



CEN-CENELEC Management Centre:
Rue de la Science 23, B-1040 Brussels

European foreword

This document (EN ISO/IEC 12792:2025) has been prepared by Technical Committee ISO/IEC JTC 1 "Information technology" in collaboration with Technical Committee CEN-CENELEC/ JTC 21 "Artificial Intelligence" the secretariat of which is held by DS.

This European Standard shall be given the status of a national standard, either by publication of an identical text or by endorsement, at the latest by May 2026, and conflicting national standards shall be withdrawn at the latest by May 2026.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. CEN-CENELEC shall not be held responsible for identifying any or all such patent rights.

Any feedback and questions on this document should be directed to the users' national standards body/national committee. A complete listing of these bodies can be found on the CEN and CENELEC websites.

According to the CEN-CENELEC Internal Regulations, the national standards organizations of the following countries are bound to implement this European Standard: Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Norway, Poland, Portugal, Republic of North Macedonia, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Türkiye and the United Kingdom.

Endorsement notice

The text of ISO/IEC 12792:2025 has been approved by CEN-CENELEC as EN ISO/IEC 12792:2025 without any modification.

Contents

Page

Foreword	v
Introduction	vi
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Symbols and abbreviated terms	4
5 Overview	4
5.1 General	4
5.2 Organization and usage of the taxonomy	4
5.3 Constraints on transparency disclosures	5
5.4 Concept of transparency	6
6 Stakeholders' needs and transparency objectives	7
6.1 General	7
6.2 Transparency objectives and goals	8
6.3 Selected stakeholder roles in transparency	9
7 Context-level taxonomy	11
7.1 General	11
7.2 Societal context	11
7.2.1 General	11
7.2.2 Labour practices	14
7.2.3 Consumer needs	15
7.3 Environmental context	16
7.3.1 General	16
7.3.2 Particular environmental disclosures	17
7.3.3 Organizational disclosures	18
8 System-level taxonomy	19
8.1 General	19
8.2 Basic information	19
8.3 Organizational processes	20
8.3.1 General	20
8.3.2 Governance	20
8.3.3 Management system	21
8.3.4 Risk management	21
8.3.5 Quality management	21
8.4 Applicability	21
8.4.1 General	21
8.4.2 Intended purposes	22
8.4.3 Capabilities	22
8.4.4 Functional limitations	22
8.4.5 Recommended uses	22
8.4.6 Precluded uses	22
8.5 Overview of technical characteristics	22
8.5.1 General	22
8.5.2 Expected inputs and outputs	23
8.5.3 Production data	23
8.5.4 Logging and storing	23
8.5.5 System decomposition	23
8.5.6 Application programming interface	24
8.5.7 Human factors	24
8.5.8 Deployment methods	24
8.5.9 Configuration management	24
8.6 Access to internal elements	25

8.7	Quality and performance	26
8.7.1	General	26
8.7.2	Verification and validation processes	26
8.7.3	Runtime measurements	26
8.7.4	Comparison with alternative systems	27
9	Model-level taxonomy	27
9.1	General	27
9.2	Basic information	27
9.3	Usage and model interplay	28
9.3.1	Processing performed by the model	28
9.3.2	Dependence on other models	28
9.3.3	Coherence with AI system's intended purposes	28
9.4	Technical characteristics	28
9.4.1	Type of technology used	28
9.4.2	Features extracted from input data	28
9.4.3	Algorithm used for processing	28
9.4.4	Procedure for building the model	28
9.4.5	Hyperparameters	29
9.4.6	Input and output formats	30
9.4.7	Compute hardware	30
9.4.8	Computational costs	31
9.4.9	Models in evolvable systems	31
9.5	Data used	31
9.6	Functional correctness	32
10	Dataset-level taxonomy	32
10.1	General	32
10.2	Basic information	33
10.3	Data provenance	33
10.4	Data properties	35
10.5	Dataset domain and purposes	36
10.5.1	General	36
10.5.2	Language data domain details	36
10.5.3	Vision data domain details	37
10.6	Data biases and limitations	37
10.7	Societal considerations	37
10.8	Data preparation performed	38
10.9	Dataset maintenance	39
	Annex A (informative) Examples of transparency templates	40
	Annex B (informative) Examples of stakeholder roles in transparency	41
	Bibliography	43

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and <https://patents.iec.ch>. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 42, *Artificial intelligence*, in collaboration with the European Committee for Standardization (CEN) Technical Committee CEN/CLC/JTC 21, *Artificial Intelligence*, in accordance with the Agreement on technical cooperation between ISO and CEN (Vienna Agreement).

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Introduction

The objectives of this document include:

- improving trustworthiness, accountability and communication among different AI stakeholders, including partners in a supply chain, customers, users, society and regulators, by establishing a consistent terminology around transparency of AI systems;
- providing AI stakeholders with information about different elements of transparency with their relevance and possible limitations to different use cases and target audience;
- serving as a basis for developing technology-specific, industry-specific or region-specific standards for transparency of AI systems.

Transparency for AI systems is the property of a system that stakeholders receive relevant information about the system. This can include information on items such as system features, limitations, data, system design and design choices (see ISO/IEC 22989:2022, 5.15.8).

Transparency enables relevant stakeholders to have access to information, so they can better understand how an AI system is developed, deployed and can be used. For example, this allows an AI customer (such as an AI user) of an AI system to determine if it is appropriate for their situation and supports an AI auditor in assessing if the system complies with conformity requirements.

A standardized transparency taxonomy of AI systems helps people with different backgrounds to better understand each other by using the same terminology. This in turn supports an improved understanding of AI systems, and provides a foundation for developing interoperable and coherent transparency related standards.

This document is structured as follows:

- [Clause 5](#) provides an overview of this document and describes the concept of transparency of AI systems;
- [Clause 6](#) discusses how transparency needs can vary depending on the AI system context and on the stakeholders involved;
- [Clause 7](#) discusses transparency items that describe the context of the AI system;
- [Clause 8](#) pertains to describing transparency disclosures at the AI system level;
- [Clause 9](#) focuses on documenting the internal functioning of the system;
- [Clause 10](#) offers guidance on the documentation of datasets as stand-alone items.

AI systems often affect society and the environment (see [Clause 7](#)). However, it is also true that society and the environment can affect the performance of an AI system. This can include various aspects, such as:

- environmental conditions that affect the functioning and longevity of an AI system;
- introduction or reinforcement of unwanted bias;
- organizational practices that can cause poor outcomes;
- effects such as concept drift (concept drift occurs when there is a change in the functional relationship between the model's input and output) and data drift (data drift is change in the statistical properties and characteristics of the input data);
- formation of unwanted feedback loops (e.g. negative reinforcement of discriminatory patterns), which can be especially problematic in the case of continuous learning.

While these are important items for consideration, a detailed analysis is not provided in this document. Some of these aspects are further detailed in ISO/IEC 23894, ISO/IEC 27701, and ISO/IEC 29134.

Information technology — Artificial intelligence (AI) — Transparency taxonomy of AI systems

1 Scope

This document specifies a taxonomy of information elements to assist AI stakeholders with identifying and addressing the needs for transparency of AI systems. The document describes the semantics of the information elements and their relevance to the various objectives of different stakeholders.

This document is applicable to any kind of organization and application involving an AI system.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 22989:2022, *Information technology — Artificial intelligence — Artificial intelligence concepts and terminology*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 22989 and the following apply.

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <https://www.electropedia.org/>

3.1

attestation

issue of a statement, based on a decision, that fulfilment of *specified requirements* (3.14) has been demonstrated

Note 1 to entry: The resulting statement is intended to convey the assurance that the specified requirements have been fulfilled. Such an assurance does not, of itself, provide contractual or other legal guarantees.

Note 2 to entry: First-party attestation and third-party attestation are distinguished by the terms declaration, certification and accreditation, but there is no corresponding term applicable to second-party attestation.

[SOURCE: ISO/IEC 17000:2020, 7.3; modified — removed 'referred to in this document as a "statement of conformity"' from Note 1.]

3.2

carbon footprint

carbon footprint of a product

CFP

sum of greenhouse gas emissions and greenhouse gas removals in a product system, expressed as carbon dioxide equivalents and based on a life cycle assessment using the single impact category of climate change

[SOURCE: ISO 14050:2020, 3.11.1; modified — added the preferred term "carbon footprint"]

3.3

confidential data

data to which only a limited number of persons have access and which are meant for restricted use

[SOURCE: ISO 5127:2017, 3.1.10.18; modified — removed Note 1 to entry]

3.4

conformity assessment body

body that performs conformity assessment activities, excluding accreditation

[SOURCE: ISO/IEC 17000:2020, 4.6]

3.5

data breach

compromise of security that leads to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to protected data transmitted, stored, or otherwise processed

[SOURCE: ISO/IEC 27040:2024, 3.5.2]

3.6

data protection

implementation of administrative, technical, or physical measures to guard against the unauthorized access to data

[SOURCE: ISO/IEC 2382:2015, 2126371; modified — removed all Notes to entry.]

3.7

digital fingerprint

bit sequence generated from a digital document using an algorithm that uniquely identifies the original document

Note 1 to entry: Any digital document modification will produce a different fingerprint.

[SOURCE: ISO 14641:2018, 3.15]

3.8

evolutive system

AI system whose behaviour can change without an explicit action from its AI developer

Note 1 to entry: Examples of evolutive systems include AI systems with a storage component, AI systems that leverage user-side actions and feedback, and AI systems involving continuous learning.

Note 2 to entry: Change of behaviour does not imply a change of goal but can result in a better or worse fulfilment of the targeted goal.

3.9

feature

measurable property of an object or event with respect to a set of characteristics

Note 1 to entry: Features play a role in training and prediction.

Note 2 to entry: Features provide a machine-readable way to describe the relevant objects. As the algorithm will not go back to the objects or events themselves, feature representations are designed to contain all useful information.

[SOURCE: ISO/IEC 23053:2022, 3.3.3; modified — removed <machine learning> domain tag]

3.10

life cycle impact assessment

LCIA

environmental impact assessment

phase of life cycle assessment aimed at understanding and evaluating the magnitude and significance of the potential environmental impacts for a product system throughout the life cycle of the product

[SOURCE: ISO 14050:2020, 3.6.5; modified — added admitted term "environmental impact assessment"]

3.11

object of conformity assessment

object

entity to which *specified requirements* ([3.14](#)) apply

EXAMPLE Product, process, service, system, installation, project, data, design, material, claim, person, body or organization, or any combination thereof.

Note 1 to entry: The term "body" is used in this definition to refer to conformity assessment bodies and accreditation bodies. The term "organization" is used in its general meaning and may include bodies according to the context. The more specific ISO/IEC Guide 2 definition of an organization as a body based on membership is not applicable to the field of conformity assessment.

[SOURCE: ISO/IEC 17000:2020, 4.2; modified — "in this document" to "in this definition" in Note 1 to entry]

3.12

power usage effectiveness

PUE

ratio of the data centre total energy consumption to information technology equipment energy consumption, calculated, measured or assessed across the same period

Note 1 to entry: Sometimes the inverse value of PUE, referred to as Data Centre Infrastructure Efficiency (DCiE), is used.

[SOURCE: ISO/IEC 30134-2:2016, 3.1.3]

3.13

scope of attestation

range or characteristics of *objects of conformity assessment* ([3.11](#)) covered by *attestation* ([3.1](#))

[SOURCE: ISO/IEC 17000:2020, 7.4]

3.14

specified requirement

need or expectation that is stated

Note 1 to entry: Specified requirements can be stated in normative documents such as regulations, standards and technical specifications.

Note 2 to entry: Specified requirements can be detailed or general.

[SOURCE: ISO/IEC 17000:2020, 5.1]

3.15

taxonomy

systematic classification of items into generic groups based on factors possibly common to several of the items

[SOURCE: ISO 14224:2016, 3.92]

3.16

water footprint

metric(s) that quantifies the potential environmental impacts related to water

[SOURCE: ISO 14050:2020, 3.10.1]

4 Symbols and abbreviated terms

AI	artificial intelligence
ML	machine learning
API	application programming interface
URL	universal resource locator
REST	representational state transfer
JSON	JavaScript object notation
GUI	graphical user interface
YAML	YAML ain't markup language

5 Overview

5.1 General

[Clause 5](#) provides an overview of this document and describes the concept of transparency of AI systems, in particular, in relation to other trustworthiness characteristics.

5.2 Organization and usage of the taxonomy

This document describes a multi-faceted taxonomy of transparency for AI systems. [Clause 6](#) discusses how transparency needs can vary, depending on the AI system context and the role of different AI stakeholders. The remainder of this document covers multiple levels and viewpoints on the system. It is organized in four parts:

- [Clause 7](#) describes transparency taxonomy elements (referred to simply as "taxonomy elements") informing on the context of the AI system (e.g. transparency needs in the case of labour stakeholders).
- [Clause 8](#) pertains to describing transparency disclosures at the AI system level (e.g. its recommended and precluded uses, or API documentation).
- [Clause 9](#) focuses on the internal functioning of the system, at the level of model components and algorithmic processes.
- [Clause 10](#) offers guidance on the documentation of datasets as stand-alone items, considering that the AI system's life cycle involves those datasets (thereby also worth transparency considerations), but also that datasets can be reused as is across several AI systems for different use cases.

For a given AI system, all four parts of the taxonomy ([Clauses 7](#) to [10](#)) can be considered, as they provide complementary information. An illustration of this document's structure can be found in [Figure 1](#).




Figure 1 — Overview of the document structure

Since the development of an AI system can involve different stakeholders (data provider, AI developer, AI platform provider, etc.) and since an AI system can be integrated into a larger AI system in a modular way, this taxonomy is also designed to be modular. In particular, the documentation of an AI system can include references to the documentation (entire, or only specific parts) of another AI system.

For example, in the case of general-purpose models released as open-source assets, it can be useful to document the models according to [Clause 9](#) but also [Clause 8](#), in order to cover recommended and precluded uses. Then another organization, when integrating the general-purpose model as an internal component of a more complex AI system, can produce a new documentation for that AI system, with:

- its own documentation according to [Clause 8](#) (e.g. recommended and precluded uses of the whole AI system);
- a reference to the general-purpose model's documentation according to [Clause 9](#);
- a new section, documented according to [Clause 9](#), for any custom model developed for that AI system;
- but discarding the [Clause 8](#) information on the general-purpose model if deemed irrelevant in the context of that particular AI system. An example of irrelevant information in this context is the recommended or precluded uses of the general-purpose model, when it includes considerations pertaining to alternate uses that are beyond the one use that is made of the general-purpose model within that AI system.

5.3 Constraints on transparency disclosures

Taxonomy elements described in all four parts of this document ([Clauses 7 to 10](#)) are elements that are useful to consider in at least some cases, but do not necessarily apply to all use cases and systems. In particular:

- There can be taxonomy elements for which disclosure is not appropriate in a given use case. For example, releasing the source code is not appropriate if it includes proprietary code ([8.6](#)). Indication of who funded a dataset's creation can be impossible if the dataset has been created by volunteers or if disclosure is prevented by contractual obligations ([10.3](#)). Similarly, revealing how much data labelers were remunerated can run counter to privacy concerns ([10.8](#)).
- There can be taxonomy elements that only exist for a given type of system. For example, information on the language variety is not necessarily applicable to a computer vision system ([10.5](#)).

- There can be taxonomy elements that describe processes that are not possible or not appropriate in a given use case. For example, comparative evaluation with a similar system can be impossible when creating a new type of AI system that has no existing alternative (8.7.4).
- Similarly, considerations for diversity and inclusion are not necessarily appropriate when producing a dataset to be representative of one specific part of the population, such as elderly females with a speech disorder (10.6).
- To determine the usefulness of transparency elements, an organization can also consider the costs and benefits associated with each of the relevant transparency elements, in order to manage potential trade-offs and define the appropriate level of transparency needed for its specific use case. For example, an organization can consider whether transparency elements can in some cases cause more harm than benefits to individuals and societies, and whether disclosures can be in breach of applicable laws and regulations or contractual requirements.

5.4 Concept of transparency

Transparency is a broad concept, but within this document, the focus lies primarily on transparency as the property of a system and, to a lesser degree, on transparency as the property of an organization. Transparency for an AI system means that appropriate information about the system is made available to relevant stakeholders (ISO/IEC 22989:2022, 5.15.8). In contrast, the transparency of an organization (referred to as organizational transparency in this document) is related to how appropriate activities and decisions are communicated to relevant stakeholders (ISO/IEC TR 24028:2020, 10.3.1).

Organizational transparency is relevant in so far as organizational principles and processes affect an AI system at various stages in its life cycle as described in ISO/IEC 22989:2022, Clause 6.

Some examples of appropriate information for transparency disclosures are: goals, known limitations, design choices and assumptions, features, models, algorithms, training methods, details of used data and quality assurance processes (for the AI system, for data used in training or evaluation). However, transparency needs can be different for different stakeholders and vary from one use case to another. Therefore, appropriate information to make available as part of AI system transparency is determined not only by who the target stakeholders are, but also depends on the use case. For example, an AI auditor or regulator can have transparency needs that differ in aspect or level of detail from the transparency needs of an AI user. For some use cases, an audit can be unnecessary; the transparency needs of individuals and societies are different when using AI systems to recommend a movie, to predict turbine repairs or to identify criminals as part of a police investigation. Additionally, for a given taxonomy element, such as the model used, transparency can be achieved for some stakeholders by sharing the documentation on the choice of model architecture used, while for other stakeholders, sharing internal states of the model, such as model parameters, can be appropriate.

The degree to which the information is useful, meaningful, truthful, comprehensive, accessible and understandable is also relevant to whether the information is appropriate. This varies according to the use cases and the type of stakeholders. For example, a non-expert user does not need the same kind of understandable information as a technical forensic investigator. The degree to which a stakeholder finds certain information accessible, especially when requesting access to the data, the code or the models, can potentially also run counter to security, privacy or confidentiality requirements. In such cases, to consider whether the transparency information can be provided to stakeholders downstream in the AI system life cycle while protecting sensitive information about upstream components, it can be necessary to secure, protect or use the information in accordance with internal organizational policies for handling of intellectual property rights, contractual obligation, confidential information or trade secrets, to identify any trade-offs.

Beyond the definitions used in this document, the concept of transparency can also be viewed as a characterization of the presentation of information, or as an openness and willingness to communicate about actions and decisions. See [Figure 2](#) for a high-level overview of the various core concepts of transparency and the relevant International Standards.




Figure 2 — Relevant high-level concepts of transparency

Transparency overlaps with, and in many cases is a prerequisite for, other AI system properties, including reproducibility, observability, verifiability, usability, explainability and interpretability, accountability, quality and trustworthiness (e.g. ISO/IEC 27036-1, ISO/IEC 27037, ISO/IEC 25059, ISO/IEC TR 24028).

For example, when ethical and societal concerns are to be considered with the use of AI systems, transparency and explainability are key for the acceptance of AI systems. Transparency in an AI system enables AI stakeholders to see how the AI system works, evaluate its functionality and comprehend its strengths and limitations. This helps develop AI systems that benefit society and address ethical and societal concerns (see ISO/IEC TR 24368).

From a software quality perspective, transparency is a sub-characteristic of usability in terms of the software quality model as defined in ISO/IEC 25059. It helps potential users to choose an AI system fitting their requirements, improving relevant stakeholders' knowledge about the applicability and the limitations of an AI system, as well as assisting with its explainability. Transparency is also a sub-characteristic of satisfaction in the quality-in-use model (see ISO/IEC 25059).

From an explainability perspective, transparency is complementary and useful to explainability. For example, some explainability methods require certain degrees of transparency to access inner states or intermediate results from an AI system, or to acquire information on the models used in AI systems (see ISO/IEC TS 6254).

6 Stakeholders' needs and transparency objectives

6.1 General

Transparency is a contributing characteristic to the trustworthiness of an AI system. Establishing and maintaining the trustworthiness of an AI system involves considering organizations, as well as individuals and groups of individuals as stakeholders. Such considerations include stakeholders that do not necessarily have a priori needs or expectations of the system. Those needs and expectations are likely to be triggered by a transparency disclosure (e.g. data subjects who are not aware that their data are being used, until informed). As is the case for other AI system trustworthiness characteristics, the benefits of transparency can be improved by being clearly contextualized, and by clarifying which stakeholders aim to establish and maintain the trust of which other stakeholders in relation to specific aspects of AI system development and use (see ISO/IEC TR 24028:2020, 8.1.1).

Transparency involves the communication of appropriate information to stakeholders in a comprehensive, accessible and understandable manner. Transparency can refer to communication of appropriate information about an AI system to relevant stakeholders or the communication of information about appropriate organizational activities and decisions to relevant stakeholders.

Transparency can involve communication of information about design goals, intended uses, known limitations, definitions, design choices and assumptions, features, components, procedures, measures, models, algorithms, training methods, quality assurance processes (for the AI system, for data used in training or evaluation) and the characteristics and provenance of data used in training and evaluation. This can also include information on the nature of processing and automation related to decisions relevant to stakeholders [ISO/IEC 22989:2022, 5.15.8].

AI stakeholder roles and sub-roles are defined in ISO/IEC 22989:2022, 5.19, and Figure 2 therein. Different stakeholder objectives influence the type of information that can be useful; for example, objectives for sustainability, which can be supported by carbon footprint analyses (see 7.3), differ from purchasing or procurement objectives, which can benefit from information on costs for usage or development (see 9.4.7).

Addressing stakeholders' needs and objectives also means tailoring the language and the level of detail of transparency information. For example, stakeholders such as certain AI user profiles or AI subjects need clear, plain language information describing the overall functioning and effect of an AI system (e.g. as described in 8.2 to 8.4); by contrast, development or deployment roles, like AI producer or AI system integrator, typically require greater levels of detail, for example, as provided in 8.5 to 8.7. ISO 24495-1 provides pertinent guidelines on how to achieve plain language, which can be relevant to achieving transparency.

Different templates for presenting information to stakeholders also exist (see Annex A for some examples currently in use in academic research and industry).

ISO/IEC 22989:2022, 5.19, identifies stakeholders and categorizes them according to their roles and sub-roles in the AI system life cycle.

A transparency relationship can have two roles as a stakeholder pair: one that provides the content and the other that is receiving the content. In order to demonstrate the various transparency responsibilities and requirements different stakeholders can have, 6.3 provides a list containing, for each stakeholder pair:

- the transparency objective;
- selected transparency items;
- a reference to the clause in which the transparency item is discussed.

The information presented in 6.3 is intended as examples only, to illustrate:

- how AI stakeholder needs can vary;
- the part different stakeholders play in providing the desired taxonomy elements.

The examples are not exhaustive and are not intended to cover all stakeholder roles and all possible stakeholder needs, nor do they cover all aspects that can also vary with the particular AI system intended use and context.

6.2 Transparency objectives and goals

Examples of transparency objectives and goals include:

- right to know: this means to be informed that humans are interacting with AI systems, and to be properly informed about the content of AI synthesis, the capabilities and limitations of the AI system, and the rights of the people affected. It can include autonomous decisions, AI system interactions, AI generated contents, and legal compliance.
- proper usage: this helps users understand the obligations, correct usage and excluded usage that apply to AI systems, thereby enabling users to use AI systems appropriately. It can include legal usage, correct usage, misuse prevention and risk mitigation.
- traceability: this transparency objective is handled across the entire AI system life cycle (design and development, verification and validation, deployment, ongoing monitoring, etc.) to help stakeholders better manage and use the AI system.

Note that the transparency objectives described here are not mutually exclusive (i.e. it is possible that a taxonomy element simultaneously satisfies multiple transparency objectives).

6.3 Selected stakeholder roles in transparency

Examples of possible transparency responsibilities and requirements pertaining to different stakeholder pairs include:

- AI provider to AI deployer:
 - for the transparency objective of right to know, the transparency contents can include:
 - basic information of system [\[8.2\]](#);
 - for the transparency objective of proper usage, the transparency contents can include:
 - applicability [\[8.4\]](#);
 - overview of technical characteristics [\[8.5\]](#);
 - access to internal elements [\[8.6\]](#);
 - for the transparency objective of traceability, the transparency contents can include:
 - basic information [\[8.2\]](#);
 - quality and performance [\[8.7\]](#);
- AI provider to AI auditor:
 - for the transparency objective of traceability, the transparency contents can include:
 - access to relevant internal elements [\[8.6\]](#);
 - the location of data processing and storage [\[8.2, 8.5.5\]](#);
- AI provider to relevant authorities:
 - for the transparency objective of traceability, the transparency contents can include:
 - compliance with legal requirements or conformity to International Standards [\[8.2\]](#);
 - the location of data processing and storage [\[8.2, 8.5.5\]](#);
- AI provider to AI end user, AI auditors:
 - for the transparency objective of traceability, the transparency contents can include:
 - dataset used considerations [\[9.5\]](#), see also ISO/IEC 42001:2023, A.7;
- AI deployer to AI end user:
 - for the transparency objective of right to know, the transparency contents can include:
 - basic system information [\[8.2\]](#);
 - for the transparency objective of proper usage, the transparency contents can include:
 - applicability [\[8.4\]](#);
 - overview of technical characteristics [\[8.5\]](#);