

Transparency of AI-generated Content: New Obligations from August 2, 2026

August 17, 2026 | WAHL Legal Alerts | ZS

As of 2 August 2026, the transparency obligations under Article 50 of Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence, known as the Artificial Intelligence Act or AI Act, become applicable. Article 50 includes four sets of obligations: (i) informing natural persons that they are interacting with an AI system, (ii) marking synthetic content in a machine-readable format, (iii) disclosing deep fake content and certain AI-generated texts, and (iv) informing natural persons exposed to emotion recognition or biometric categorisation systems.

These obligations may be relevant for companies that use off-the-shelf AI tools, for example in customer support, marketing, public relations, and other business processes. Every organisation should thus determine how it uses AI and whether, in a given case, it acts as a provider or as a deployer of an AI system.

I. WHO MUST COMPLY?

The AI Act distinguishes between two roles:

- (i) **The provider of an AI system** is an entity that develops an AI system, or places it on the market or puts it into service under its own name or trademark.
- (ii) **The deployer of an AI system** is a natural or legal person that uses an AI system in the course of its professional activity.

For example, a company that develops a chatbot and offers it to other business users will generally be considered the provider of the AI system. A company that deploys that chatbot on its own website to communicate with customers will usually be the deployer of the AI system.

Providers and deployers of AI systems **do not have the same obligations**. Providers are primarily responsible for the design of the system and for the technical marking of its outputs, whereas deployers are responsible for clearly informing natural persons and for visibly labelling certain types of content.

II. KEY TRANSPARENCY OBLIGATIONS

Article 50 of the AI Act lays down transparency obligations that depend on the type of AI system, its use, and the content it generates or modifies. The key obligations are set out below.

Users must know when they are interacting with AI

Providers of AI systems intended to interact directly with natural persons must ensure that users are informed that they are interacting with AI.

This obligation will be particularly important for chatbots, virtual assistants, automated voice systems, and other systems that communicate with users autonomously.

The notice must be given at the latest at the time of the first interaction. It also must be given in a clear and distinguishable manner, and must not remain hidden solely within a privacy policy or general terms and conditions. In practice, a simple notice such as the following may be used:

"You are talking to a virtual assistant that uses artificial intelligence."

A notice is not required where, taking into account the circumstances and the context of use, it is obvious to a reasonably well-informed, observant and circumspect person that they are interacting with an AI system. Since the assessment of obviousness may depend on the specific case, it is advisable to clearly inform users whenever there is any doubt.

A special exception to this rule exists for certain systems authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to appropriate safeguards. This exception does not apply where such systems are available to the public for reporting criminal offences.

Technical marking of AI-generated content

Providers of AI systems, including general purpose AI systems, that generate synthetic content such as audio, image, video or text content must ensure that the output is marked in a machine-readable format and is technically detectable as artificially generated or manipulated.

A machine-readable marking is not the same as a visible notice stating that content was *"created with the help of AI."* It is a technical marking that can be recognised by other systems used to verify the content provenance.

The AI Act requires the technical solutions used for marking to be effective, interoperable, robust and reliable as far as this is technically feasible, taking into account the specificities and limitations of the various types of content, the

costs of implementation, and the generally acknowledged state of the art.

Since a single marking method is often not sufficient, a layered approach may be applied in practice. For example, digitally signed metadata can be combined with an imperceptible digital watermark.

Digitally signed metadata can contain information on whether the content was generated or modified using AI. An imperceptible watermark is embedded within the content itself so that it remains difficult to remove even after the content has been altered or shared.

The obligation to apply machine-readable marking does not apply where the AI system performs solely an assistive function for standard editing, or does not substantially alter the input content or its meaning. An exception also exists for certain systems authorised by law to detect, prevent, investigate or prosecute criminal offences.

Preservation and verification of technical markings

For technical solutions to remain effective, robust and reliable, providers should make reasonable efforts to preserve existing technical markings on content that is fed into their system and subsequently processed further. It is advisable to prohibit, in the terms of use or other documents, intentionally removing or altering markings, as well as offering or promoting of tools designed to circumvent such markings.

Output content must be technically recognisable as artificially generated or manipulated, which is why providers should ensure an appropriate solution for verifying its provenance.

The detection result must be presented in a clear and understandable manner. It should also, where possible, explain the technique on which it is based.

It is important to emphasise that the detection result indicates the possible provenance of the content but does not determine its truthfulness. The fact that content was generated using artificial intelligence does not necessarily mean

that it is inaccurate. Equally, the absence of a technical marking does not prove that the content is authentic.

Where a user is required to upload content into a detection tool, the provider of that tool must comply with the rules on the protection of personal data and confidentiality. Uploaded content should be used only for the purpose of verification and deleted once the verification has been completed. The exception exists for a limited set of technical data whose retention is necessary for security purposes or to prevent misuse.

Deepfake content must be clearly labeled

A deployer that uses an AI system to generate or manipulate content constituting a deep fake must clearly disclose that the content has been artificially generated or manipulated.

A deep fake is image, audio or video content that has been generated or manipulated using AI and that resembles existing persons, objects, places, organisations or events in a way that would falsely appear to be authentic or truthful. For example, this may be a video that makes it appear as though a real person is saying something they never said, or an audio recording in which a real person's voice appears to say something that person never said (voice cloning).

The label must be clear and noticeable at the latest when the person is first exposed to the content; it cannot exist solely in technical metadata.

The visible label applied by the deployer is different from the machine-readable marking for which the provider of the AI system is responsible, so in certain cases both obligations may apply at the same time.

The following wording may be used:

„This content was generated using artificial intelligence.“

Or

„This content has been modified using artificial intelligence.“

Standardised icons or other clear textual and visual markings may be used for labelling purposes. For images or videos, the label should be placed sufficiently close to the content. For videos and audio recordings it is advisable to display the label at the beginning and, as far along as possible, to retain it when the content is shared.

III. RULES FOR SPECIFIC PURPOSES AND USES

Certain types of content and ways of using AI systems require a specific approach to transparency. This applies in particular to artistic and creative works, texts on matters of public interest, and emotion recognition and biometric categorisation systems.

Artistic, creative and satirical works

Where a deep fake forms part of an evidently artistic, creative, satirical, fictional or similar work or programme, the transparency obligation continues to apply, but the manner of labelling may be adapted so as not to unnecessarily hamper the display or enjoyment of the work.

For example, in a film where AI has been used to de-age an actor, it is not necessary to display a large label over the image throughout the entire film. The information may instead be stated in the description of the work, in the opening or closing credits, or in another appropriate place.

AI-generated texts on matters of public interest

A specific obligation applies to text that has been generated or modified using an AI system and published for the purpose of informing the public on matters of public interest.

Matters of public interest may include politics, elections, healthcare, public safety, the economy, the environment, education, the judiciary, the activities of public bodies, and other topics of importance to a large number of citizens.

Such text must be labelled as AI-generated or modified, unless two conditions are met at the same time: (1) the content has undergone

human review or editorial control, and (2) a natural or legal person has assumed editorial responsibility for its publication.

Human review should not be merely formal, but should include verification of the facts, sources, and compliance with applicable regulations and internal rules. The use of AI tools solely as an aid in preparing a first draft does not automatically trigger the labelling obligation, provided that the final content is genuinely reviewed and there is clear editorial responsibility for its publication.

Where no such control exists, the label should be placed above or near the top of the text, close to the title, or in another immediately visible location.

Emotion recognition and biometric categorisation

Deployers of emotion recognition or biometric categorisation systems must inform the natural persons exposed to such systems of their operation. These are systems that attempt, on the basis of biometric data, to infer a person's emotional state or to classify them into a particular category. Such systems may be used, for example, in customer support, vehicles, security systems, retail, or visitor analytics.

The obligation to inform does not mean that every use of such systems is permitted. For certain applications, the AI Act provides for prohibitions or stricter conditions, and the GDPR and other privacy rules must be complied with when processing personal data. Information on the operation of the system must be provided in a clear and distinguishable manner at the latest when the person is first exposed to the system, and must meet the applicable accessibility requirements.

IV. QUALITY OF TECHNICAL SOLUTIONS

Article 50(2) of the AI Act lays down four basic requirements for technical solutions: effectiveness, interoperability, robustness and reliability.

Effectiveness means that the technical marking enables the recognition of artificially generated or manipulated content. Reliability requires that the solution distinguishes marked AI content from other content with sufficient accuracy. Robustness means that the marking, as far as this is technically feasible, remains preserved following ordinary or intentional alterations to the content. Interoperability mandates that the markings and their verification can function across different systems, platforms and technical environments.

For the practical implementation of the interoperability requirement, it is advisable to use appropriate technical standards and solutions that enable different marking and detection systems to be interconnected.

V. INTERNAL PROCEDURES AND DOCUMENTATION

Organisations should establish internal procedures that specify when content must be labelled, which label is to be used, where it is to be placed, who is responsible for verifying it, and how defective or missing labels are to be corrected.

Employees and external collaborators who create or publish content should be familiar with these internal procedures. Providers of generative AI systems should document the technical solutions they use, how those solutions are tested, any shortcomings identified, and the corrective measures taken. Solutions should be tested before the system is placed on the market or put into service, and regularly throughout its life cycle, under conditions that correspond to actual use.

The scope of documentation, testing and training may be adapted to the organisation's role, the type of AI system, and its size and available resources, provided that the measures taken genuinely enable compliance with the applicable obligations under AI Act Article 50.

VI. RELATION TO OTHER REGULATIONS

The obligations under Article 50 of the AI Act do not apply in isolation from other rules and applicable regulations.

Where an AI system processes personal data, the GDPR applies. This is particularly important in the case of deep fakes depicting real persons, voice cloning, emotion recognition, biometric categorisation, and tools to which users upload content to verify AI markings.

Depending on the content type and the sector concerned, rules on consumer protection, intellectual property, digital services, media, political advertising, protection of personality rights, criminal law and accessibility may also apply.

A clear label does not remove the need to obtain the rights required to use protected content. Likewise, properly labelling deep fake content does not mean that its publication cannot amount to an infringement of personality rights, misleading advertising, or other forms of unlawful conduct.

VII. PENALTIES

A breach of the transparency obligations under Article 50 of the AI Act may result in administrative fines of up to EUR 15,000,000 or, if the offender is an undertaking, up to 3% of its total worldwide annual turnover for the preceding financial year. Large enterprises are subject to whichever of the two amounts is higher, whereas for small and medium-sized enterprises (SMEs), including start-ups, whichever of the two amounts applies.

When determining the fine, account is taken of the nature, gravity and duration of the infringement, whether it was intentional or negligent, the mitigating measures taken, cooperation with the competent authorities, previous infringements, and the size of the enterprise.

In addition to fines, the competent authorities may also require other corrective measures.

VIII. WHAT SHOULD ORGANISATIONS DO?

1. All AI systems that are developed, procured or used should be recorded. For each system, its purpose, the types of content it generates or modifies, the persons exposed to it, and the organisation's role under the AI Act should be identified.
2. Particular attention should be paid to the use of chatbots and virtual assistants, the generation of AI content, the publication of content that could constitute a deep fake, the preparation of texts on matters of public interest, and the use of emotion recognition or biometric categorisation systems.
3. The organisation should review the technical capabilities of its AI tool providers and its contractual relationships with them. In doing so, it is important to establish whether machine-readable markings exist, whether they are retained after the content is downloaded or edited, whether a tool exists for verifying them, and what happens to the data submitted for verification.
4. Organisations that publish AI-generated content should establish internal rules on visible labelling. These rules should specify which wording or icons are used, where they are placed, who verifies them before publication, and what procedure applies if a label is lost through further sharing.
5. Ftexts on matters of public interest, the person or function responsible for reviewing the content and assuming responsibility for its publication must be clearly designated, and appropriate documentation of the review carried out must be ensured.
6. Employees and other persons in relevant roles should be provided with appropriate training on the applicable transparency obligations and the organisation's internal rules.

CONCLUDING REMARKS

The transparency obligations under Article 50 of the AI Act represent one of the first widely

applicable and practically visible requirements of the AI Act.

As of 2 August 2026, organisations will need to be able to demonstrate that they have complied with the obligations under Article 50 of the AI Act applicable to their AI systems and the ways in which they are used.

Given the need to combine legal, technical and organisational measures, it is advisable to begin the assessment without delay. Organisations that establish internal procedures for the transparent use of AI systems in good time will be in a significantly better position to ensure compliance and reduce regulatory and reputational risks.

Contact:

Kristina Rudec Kovačić

+385 (0)99 589 8096

k.rudec-kovacic@wahl.hr

Ivan Luetić, LL.M. Harvard Law School

+385 (0)91 588 5664

i.luetic@wahl.hr

Andrea Smolić

+385 (0)95 207 0009

a.smolic@wahl.hr

Helena Pongrac

+385 (0)95 207 0008

h.pongrac@wahl.hr

** This publication was prepared with assistance of the associate Ela Židov*

This publication was prepared by the law firm WAHL & Partners d.o.o. as a legal news alert intended for clients, associates, and partners. The information contained in this publication does not constitute legal advice and cannot be construed as such. If you have any questions or concerns regarding the content of this publication, please contact the attorney you regularly consult.