

Stakeholder consultation on the draft Guidelines on the transparency requirements for certain AI systems under Article 50 AI Act

Fields marked with * are mandatory.

Stakeholder consultation on the draft Guidelines on the transparency requirements for certain AI systems under Article 50 AI Act

Disclaimer: This document of the AI Office is prepared for the purpose of consultation and does not prejudge the final decision that the Commission may take on the final text of the guidelines on transparency requirements under Article 50 AI Act. The responses to this consultation will provide input for the guidelines on the transparency requirements for certain AI systems under Article 50 AI Act.

This consultation is targeted to stakeholders of different categories, including, but not limited to, providers and deployers of interactive and generative AI models and systems, providers and deployers of biometric categorisation and emotion recognition systems, private and public sector organisations using such interactive and generative AI systems, as well as academia and research institutions, civil society organisations, governments, supervisory authorities and the general public.

The Artificial Intelligence Act ('the AI Act'), which entered into force on 1 August 2024, creates a single market and harmonised rules for trustworthy and human-centric Artificial Intelligence (AI) in the EU. It aims to promote innovation and the uptake of AI, while ensuring a high level of protection of health, safety and fundamental rights, including democracy and the rule of law.

Among various obligations, trustworthiness of AI systems is ensured through a set of transparency obligations in Article 50 AI Act. These transparency obligations are applicable as of 2 August 2026.

They aim to enable natural persons to recognise interaction with and content generated or manipulated by AI systems, thus reducing the risks of impersonation, deception or anthropomorphisation and fostering trust and integrity in the information ecosystem.

Pursuant to Article 96(1)(d) AI Act, the Commission shall issue guidelines on the practical implementation of transparency obligations laid down in Article 50 AI Act.

The purpose of the present targeted stakeholder consultation is to collect input from a wide range of stakeholders on the draft Commission guidelines on the application of the transparency obligations in Article 50 AI Act.

The drafting of these Guidelines was informed by input from a variety of stakeholders collected during a broad consultation organised by the Commission and input from the Member States in the AI Board. The draft guidelines are now published for additional stakeholder feedback before they are formally adopted by the Commission.

The targeted consultation **is available in English only** and will be **open for 4 weeks starting on 8 May until 3 June 2026**.

All contributions to this consultation may be made publicly available. Therefore, please do not share any personal or confidential information in your contribution. It is your responsibility to avoid personal data and any reference in your contribution itself that would reveal your identity.

The guidelines will be complemented with a Code of Practice that is under development to support the practical and effective implementation of the requirements in Article 50(2) and (4) on marking and labelling of AI generated content.

This consultation accompanies the draft Guidelines on the transparency obligations under Article 50, which aim to support a consistent interpretation and effective implementation of the transparency requirements applicable to certain AI systems and AI-generated or manipulated content across the Union.

This consultation seeks to collect targeted feedback on the draft Guidelines prepared by the Commission based on broad stakeholder input. In particular, the Commission seeks feedback on whether the explanations, concepts and examples provided sufficiently support stakeholders in understanding and complying with the transparency obligations laid down in Article 50 AI Act, including when and how such obligations apply in practice. The draft Guidelines address the transparency obligations under Article 50 AI Act and are structured as follows:

- **Section I** presents an introduction outlining the background, objectives and legal context of the present Guidelines. It recalls the rationale of transparency within the AI Act, including its links to fundamental rights, user awareness and trust, and situates Article 50 within the broader risk-based framework of the AI Act.

- **Section II** provides an overview of the transparency obligations and related horizontal topics. It includes an explanation of the different obligations under Article 50, the actors responsible for their compliance, exclusions from scope (such as purely personal non-professional activities, research and development), and the interplay with other provisions of the AI Act, including prohibited practices and high-risk AI systems and general-purpose AI models and systems.
- **Section III** addresses transparency obligations for AI systems intended to interact directly with natural persons under Article 50(1) AI Act. It sets out the main concepts and scope of application, details the information obligation, and explains the relevant exceptions (including cases of obvious interaction and certain law enforcement uses), as well as the interplay with other Union legal acts.
- **Section IV** provides guidance on the marking and detection of AI-generated or manipulated content under Article 50(2) AI Act. It clarifies the scope of application, including different modalities of synthetic content, and explains the transparency obligation (marking and detection) and the technical requirements (effectiveness, reliability, robustness and interoperability). It also addresses relevant exceptions and the interaction with other Union legal frameworks.
- **Section V** covers transparency obligations for emotion recognition systems and biometric categorisation systems under Article 50(3) AI Act. It outlines the main concepts, scope and applicable obligations, as well as situations falling outside the scope and the interaction with other Union legal acts.
- **Section VI** addresses the labelling of deep fakes and AI-generated or manipulated text published to inform the public on matters of public interest under Article 50(4) AI Act. It provides clarification on the notions of e.g. deep fakes and matters of public interest, the applicable disclosure obligations (including for evidently creative, artistic, satirical, fictional or analogous works or programmes), and the relevant exceptions. It also explains the relationship with other Union legal acts.
- **Section VII** sets out horizontal requirements applicable to the information provided under Article 50(5) AI Act, including general principles ensuring that transparency information is clear, meaningful and accessible to users.

All participants are invited to provide feedback on the particular sections of the draft Guidelines they are interested in. As not all sections may be relevant for all stakeholders, respondents may reply only to the section (s) they would like. Respondents are encouraged to provide explanations and practical cases as part of their responses to support the practical usefulness of the Guidelines. We kindly ask the respondents to specify the exact section and paragraph of the draft Guidelines to which their comments refer.

The feedback collected through this consultation will support the Commission in refining and finalising the Guidelines, with the objective of ensuring that they are clear, comprehensive and practically useful for all stakeholders involved in the development, deployment and supervision of AI systems within the scope of Article 50 of the AI Act.

All contributions to this consultation may be made publicly available.

Therefore, please do not share any personal or confidential information in your contribution (your written feedback). It is your responsibility to avoid personal data and any reference in your written feedback itself that would reveal your identity.

For information on how the Commission processes your personal data please read our Privacy Statement

[7 5 26 Public Consultation Transparency50 PrivacyStatement.pdf](#)

Introductory section. Information about the respondent

* First name

* Surname

* Email address

* Do you agree that we may publish your identity together with your contribution in the instance that all contributions are made publicly available?

If you act in your personal capacity: All contributions to this consultation may be made publicly available. You can choose whether you would like your details to be made public or to remain anonymous. The respondent category that you selected for this consultation, your answer regarding residence, and your contribution may be published as received. Should you choose to remain anonymous, your name will not be published. Please do not include any personal data in the contribution itself.

If you represent one or more organisations: All contributions to this consultation may be made publicly available. You can choose whether you would like respondent details to be made public or to remain anonymous. Only the following organisation details may be published: The respondent category that you selected for this consultation, the name of the organisation on whose behalf you reply as well as its size, its presence in or outside the EU and your

contribution as received. Should you choose to remain anonymous, your name will not be published. Please do not include any personal data in the contribution itself if you want to remain anonymous.

- ☒ Yes
☐ No

* Do you agree that we may contact you in the event of follow-up questions or if we want to learn more about your responses?

- ☒ Yes
☐ No

* Do you represent an organisation (e.g., a public organisation, a company, a think tank or a civil society /consumer organisation) or act in your personal capacity (e.g., independent expert)?

- ☒ Organisation
☐ In a personal capacity

* If you are representing an organisation, please specify the name of the organisation:

Video Games Europe

* Type of organisation

Association

* Is a representation of the organisation located in the EU?

- ☒ The organisation's headquarter is located in the EU
☐ A branch office, or any representation of the organisation is located in the EU
☐ None of the representations of the organisation is located in the EU

* Select the EU Member State where the organisation's headquarter, or representation is located

BE - Belgium

* Select the size of the organisation

Micro (0-9 employees)

* Sector(s) of activity

- | | | |
|---|---|--|
| ☐ Information technology | ☐ Employment | ☐ Transport |
| ☐ Public administration | ☐ Education and training | ☐ Telecommunications |
| ☐ Law enforcement | ☐ Consumer services | ☐ Retail |
| ☐ Justice sector | ☐ Business services | ☐ E-commerce |
| ☐ Legal services sector | ☐ Banking and finances | ☐ Advertising |
| ☒ Cultural and creative sector, including media | ☐ Manufacturing | ☐ Consumer protection |
| ☐ Healthcare | ☐ Energy | ☐ Others |

* Describe the activities of your organisation or yourself

1300 character(s) maximum

Since 1998, Video Games Europe has ensured that the voice of a responsible games ecosystem is heard and understood. Its mission is to support and celebrate the sector's creative and economic potential and to ensure that players around the world enjoy the benefits of great video game playing experiences.

When providing feedback, please ensure that your responses are concise and clearly reference the relevant section, subsection and/or paragraph number(s) of the draft Guidelines.

This will support the Commission in efficiently analysing stakeholder input and refining the Guidelines to ensure they are clear, comprehensive and practically usable.

* **On which part(s) of the draft Guidelines would like comment?** *Multiple answers are possible. Please note that selecting a particular answer will direct you to a set of questions specifically related to subject specified.*

- ☒ **Section I - Background and objectives**
- ☒ **Section II - Overview of transparency obligations and horizontal topics**
- ☒ **Section III - Article 50(1): Transparency for interactive AI systems**
- ☒ **Section IV - Article 50(2): Marking and detection of AI-generated or manipulated content**
- ☐ **Section V - Article 50(3): Emotion recognition and biometric categorisation systems**
- ☒ **Section VI - Article 50(4): Labelling of deep fakes and certain text publications**
- ☒ **Section VII - Horizontal requirements (Article 50(5))**

Section I - Background and objectives

Question A. Are there any aspects / paragraphs of this Section of the draft Guidelines that you believe require further clarification?

- ☒ Yes
- ☐ No
- ☐ NA

If so, how would you suggest they be improved to ensure effective implementation and compliance?

2000 character(s) maximum

We welcome the reference to the risk-based approach of the AI Act (see §2). The European Commission itself had classified "AI-enabled video games" as an example for the "minimal or no risk" category. However, based on the contents of the draft guidelines, we are concerned that this risk classification is now contradicted by other sections of the draft guidelines which seem to impose transparency obligations on video games.

Art.50 should be interpreted consistently with its purpose, which is in addressing genuine risks of deception, impersonation, misinformation, manipulation and/or fraud at scale. Ordinary videogames are authored, branded and published entertainment products. Players understand that game worlds, characters, dialogue, animations and visual assets are simulated or fictional, even where they are visually realistic or AI-assisted.

Moreover, we believe that any disclosure requirements to be imposed should do so bearing in mind the nature of the work in question, the context in which it appears, and should remain proportionate to the goals of the AI Act. To that end, we suggest an addition to §3: "(...) in ensuring compliance with the transparency obligations under Article 50 AI Act in a consistent, *proportionate*, effective and uniform manner." Proportionality is a key element in ensuring that the application of Art.50's requirements can be done bearing in mind the context and the practical realities of a given case.

Section II - Overview of transparency obligations and horizontal topics

Question A. Are there any aspects / sections / paragraphs that require further clarification or refinement?

- ☒ Yes
☐ No

Please indicate specific parts.

2000 character(s) maximum

In §9, the purpose of the transparency obligations is linked to the risks of impersonation/deception /misinformation/manipulation at scale, while also referring to the integrity of the information ecosystem. Regarding the entertainment sphere, particularly within curated, edited and published works delivered through mainstream channels, such as commercial video games, these risks are not material, as no reasonable person expects video game imagery to correspond to actual reality. Consequently, it is not necessary to require the same disclosure obligations for situations that present minimum or no level of risk: AI-assisted imagery of a real racing circuit in a racing video game cannot be subject to similar classification as for example, insidious AI-manipulated deep fake political speeches by prominent politicians posted on social media, specifically in what concerns their characterization as a 'deep fake'. We therefore request that the final guidelines will explicitly exclude AI-enabled video games from the scope of Art.50.

For game content, any potentially applicable transparency requirements (without conceding that these apply or should apply to video game imagery and content) should be satisfied, where necessary, through proportionate, context-appropriate, and non-intrusive means, including but not limited to: end-user agreements or terms of service, credits, an AI information page within the platform/store information, parental information, game settings disclosures or game manuals. Repeated in-game labels or overlays would create habituation effects, impair gameplay and add little meaningful information for consumers. Our response to Section VII will further expand on this point.

Section III - Article 50(1): Transparency for interactive AI systems

Question A. Are there aspects requiring further clarification for effective implementation and compliance?

- ☒ Yes
- ☐ No

Please indicate specific parts.

2000 character(s) maximum

We welcome clarification regarding the extent to which the 'obvious' exception applies, particularly in the context of entertainment: does this 'obvious' exception extend to all features of a video game, including its imagery, associated marketing/promotional materials, such as trailers or teasers? Moreover, we welcome the reference to habituation effects under §36, as repeated notifications/overlays disclosing the presence of AI manipulated-edited content in synthetic mediums such as video games would not only be detrimental to consumer experience, but would also not add any meaningful information or 'user understanding' (as referenced in §34) in a context (video games) where there are no expectations of 'reality'. We also welcome §41, as it provides a standard for 'average' natural person, also considering the level of digital literacy of the target audience – which in the case of the video games industry are the players, and the standard of assessment must be consistent with the digital literacy/circumspection expected from an average consumer.

Question B. Are there additional practical examples or use-cases that should be clarified?

- ☒ Yes
- ☐ No

Please specify:

2000 character(s) maximum

The examples listed on p.11 include a reference to 'AI avatars', as a type of system that falls within the scope of Art.50, but simultaneously refers on p.15 that "interactions with AI-enabled Non-Playable Characters (NPCs) in a videogame" is obvious. While we certainly welcome the latter reference, we would expect that the former is clarified.

In the interest of legal certainty, we would like to understand where and how the draft guidelines distinguish the NPC case from the other example referenced on p.15: "AI systems embedded in immersive environments (e.g. virtual or augmented reality) using realistic avatars or voices resembling humans, where users and particularly children, elderly individuals, or persons with disabilities may not be able to distinguish easily between human and AI interaction". The 'AI avatars' example raises questions, and since §42 refers to the visibility of mechanical components as an element of obviousness, it can be argued that by the mere act of a user putting on a VR headset, calibrating the visuals, and adapting it to his/her head, should constitute sufficient indication to that user that they are engaging with a synthetic medium, and this should be provided as an example of obvious interactions on p.15.

Section IV - Article 50(2): Marking and detection of AI-generated or manipulated content

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation of marking and detection obligations?

- ☒ Yes
- ☐ No

Please specify:

2000 character(s) maximum

We welcome clarification on how Art.50(2) can effectively be implemented to dynamic, interactive, and mixed-origin environments such as online games and immersive experiences. The guidelines state that virtual, augmented, and mixed reality may be regarded as 'video' content, and that AI systems generating or manipulating 3-D images, videos, audio, and virtual environments fall within Art.50(2). The guidelines should clarify how this applies to real-time, runtime-rendered gameplay where there may be no fixed media file or stable output to mark. In particular, the Commission should clarify whether the relevant 'output' is the individual AI-generated asset, the rendered frame, the full gameplay session, or the experience as a whole. The guidelines should also provide clearer treatment of ephemeral, real-time content in video games. In §82, which we welcome, the Commission appropriately recognizes that real-time content generation that is consumed immediately and not stored or disseminated further, including in video games, may be exempt where users are aware that the content is AI-generated or manipulated. We support this approach, but the final guidelines could expand it into clearer guidance for interactive environments. We suggest the inclusion of the following language: "where AI-generated content is generated or rendered in real time, consumed during use or gameplay, not separately exported by the provider, and accompanied by reasonable user-facing transparency, machine-readable marking should not be required if watermarking is technically infeasible or provides only limited transparency value."

Question B. Are there additional examples, use-cases or technical approaches that should be included?

- ☒ Yes
☐ No

Please specify:

2000 character(s) maximum

We respectfully request that the Commission add examples addressing dynamic, interactive, and mixed-origin online game environments. In particular, the final guidelines should clarify how Art.50(2) applies where an AI-generated asset is only one component of a broader human-created or user-created experience, such as AI-generated trees, textures, terrain elements, sounds, or background objects placed within an otherwise non-AI-generated game. We believe that §60 contains positive and workable interpretational elements, and considering that, given the Commission's characterization of 'AI-enabled video games' as minimal to no risk, a reference to, for example, in-game character creation systems or game editing tools in a given game, made available for consumers, should not be subject to Art.50(2) requirements. Additionally, the Commission should also provide some examples of what they understand as real-time and ephemeral gameplay content. Following the logic of §82, the guidelines could refer to examples such as "AI-generated non-player character dialogue response, procedurally generated scenery/lighting, or AI-assisted environmental variation that appears during gameplay, which is consumed immediately, and is not stored nor separately disseminated by the provider" as types of content exempted.

Question C. What technical approaches or solutions should be considered for the implementation of the marking and detection obligations under Article 50(2) AI Act, in particular in relation to AI agents and virtual or immersive environments (e.g. virtual reality)? Please provide concrete examples, methodologies or tools where possible.

1500 character(s) maximum

Section VI - Article 50(4): Labelling of deep fakes and certain text

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation?

- ☒ Yes
☐ No

Please indicate specific parts

2000 character(s) maximum

We welcome §107i, there should be a case-by-case assessment by the deployer, based on objective comparison, in relation to what constitutes or not 'resemblance'.

However, we are concerned we the proposed extensive interpretation in §107ii, according to which a 'deep fake' includes "simulated persons, objects, places, entities or events to resemble someone or something that can exist or could have existed in reality". This goes well beyond the wording of Art. 3(60), which is limited to "existing persons, objects, places, entities or events", and should be amended accordingly.

The reference in §107ii to 'unrealistic' depictions causes ambiguity: would a highly realistic depiction of a person (for example, the user's uploaded face) in a multiplayer competitive dancing video game, engaging in extreme dance moves, be 'unrealistic'? We believe it is important to consider the 'obvious' element referred in 50(1) to be useful here, and we encourage the Commission to make use of it, to ensure that the interpretation of these provisions remains grounded in proportionality and practical user experience/expectations.

Regarding §107iv, we agree that this criterion should relate to the essential characteristic of deep fake content, namely its potential for deception or misleading: video games typically do not have such potential, and consequently, characterizing game imagery as deep fakes is fundamentally inaccurate.

There is also a contradiction with the example on p.15 (interaction with NPCs): a user who knows that he's interacting with an AI-controlled NPC will also be aware that the image of a professional football player in a football simulation game is not real. Another concerning element is that §108 collides with the notion of 'average' user laid out in §41. We suggest alignment of these two notions of 'average' in line with the context and consumer expectations in relation to the given type of content, in the interest of practicality and coherence across Art.50.

Question B. Are there additional examples or use-cases needed that should be clarified?

- ☒ Yes
☐ No

Please specify:

2000 character(s) maximum

It should be clarified that the content in video games is out of scope of Art. 50(4), as there is no element of false appearance as required by Art. 3(60). When playing a video game, the average user is clearly aware they are engaging with content apart from reality. They are not exposed to deception regarding the content, as they are aware that this is creative/fictional. Treating such curated, branded, age-rated and edited content in the same way as malicious deep fakes posted on social media would be disproportionate, and run counter the logic of the AI Act and the European Commission's simplification agenda.

We would welcome clarification regarding where assisting/editing techniques such as AI-assisted photogrammetry or Gaussian splatting, lie. Specifically, in relation to what the draft guidelines lay out in §109 for 'AI-supported minor manipulations of pre-existing content' not being deep fakes, or whether the classification of a given piece of AI-edited content as a deep fake is ultimately case/context-specific.

The examples listed on p.29 do not offer much in the way of clarity: would it mean, for video games, that any high fantasy game would benefit from a blanket exclusion from the 'deep fake' classification, whereas a racing game featuring 3D, AI-assisted recreations of real car models and race tracks, would not?

The guidelines should distinguish between AI-generated consumer-facing content and ordinary AI assisted production workflows. In games, AI may be used as part of a wider creative and technical pipeline, including concept iteration, asset optimisation, localisation, animation assistance, upscaling, and procedural generation. These uses should not trigger player facing labelling obligations where the final content is incorporated into a curated, edited, age-rated and published work. The classification as a deep fake should remain context specific, and not turn solely on whether AI was used at some stage in production.

Section VII - Horizontal requirements (Article 50(5))

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation?

Please indicate specific parts.

2000 character(s) maximum

As the artistic/creative exception is confirmed in the table in p.5 as a 'special disclosure regime', it constitutes *lex specialis* in relation also to the requirements in 50(5). Consequently, the requirement to disclose a 'deep fake' at first exposure does not apply. Rather, per §115, there is some discretion for deployers on how to inform viewers/users, and consequently, we believe that the reference on p. 37 that "Disclosure as part of end credits does not comply with Article 50(5) AI Act" is erroneous in the context of the special disclosure regime, since if the first exposure requirement were to apply to it, the exception would be functionally void, and in practice, bereft of practical utility for deployers. Thus, disclosure via credits should be a valid option for cases benefitting from the 50(4) artistic exception. It is key to recall Recital 134 of the AI Act: "(...) In those cases [50(4) artistic disclosure regime], the transparency obligation (...) is limited to disclosure (...) in an appropriate manner that does not hamper the display or enjoyment of the work, including its normal exploitation and use, while maintaining the utility and quality (...)", which clearly subordinate the manner of disclosure to the normal exploitation and use of a work, without impacting its quality.

Art.50(4) says disclosure must be made in a way not hampering display or enjoyment, whereas 50(5) says disclosure must occur no later than first exposure. In most artistic contexts, the moment of first exposure is precisely where disclosure would most hamper enjoyment. Games aren't linear works: many are persistent, procedural, adaptive, multiplayer, or user-driven, for which a rigid 'first exposure' disclosure rule is practically unworkable. We consider this reinforces the argument that Art.50(5) must be read in a practical and functional manner.

Contact

[Contact Form](#)