WARDYŃSKI & PARTNERS

SHOOTME LABS

„May contain trace amounts of AI”

Julia Padlewska-Oleksy
Monika A. Górska
Lena Marcinoska-Boulangé
Bartosz Kurzec
Krzysztof Wojdyło
Marysia Surzycka
Szymon Furi Brzozowski
Maciej Miliszkiewicz

JULY 2026

THE DUTY TO LABEL AI CONTENT

More provisions of the European Union's AI Act will soon enter into force. For the creative industry (but not only) this will bring specific new obligations to label content generated or modified with the use of artificial intelligence systems.

How to properly comply with this obligation, how to interpret the notion of "deep fakes," what should be regarded as falling within the exception of the "artistic clause," and finally, how to specifically label such materials: none of this is simple, and it raises many doubts.

This is why the law firm of Wardyński & Partners, together with the production house and creative agency SHOOTME | B Corp™ and the specialised division of SHOOTME AI Labs, organised a webinar to examine what the duty to indicate AI-generated content means in practice for content creators, producers, brands and others. Nearly 250 people took part in the webinar. This is proof enough that the topic is not merely academic. And we're talking about just one provision of the AI Act. The discussion lasted over two hours but failed to exhaust all of the questions raised.

In this joint publication, we summarise key strands of the debate and address in one place the main questions raised at the event. But we know this doesn't exhaust the topic. We will meet again after the rules on labelling of AI-generated content begin to be applied, so we can discuss together the areas still raising controversies and questions.

We would like to take this opportunity to invite you now to participate in our next event.

See you soon.

Wardyński & Partners and SHOOTME

Introduction, legal context

The AI Act (Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)) is the key law establishing the framework for use of AI systems in the EU. One of the aims of the regulation is to ensure an adequate level of transparency via the obligation to appropriately label content generated or manipulated using AI systems.

The AI Act contains several provisions touching on this obligation. We will focus on the one that is most vital for people using AI tools to create content. This is Art. 50(4) of the act, insofar as it addresses images, audio and video: "Deployers of an AI system that generates or manipulates image, audio or video content constituting a deep fake, shall disclose that the content has been artificially generated or manipulated."

This requires appropriate designation of content if:
- You are a "deployer" (as defined in the AI Act)
- You are using an "AI system" (defined in the AI Act), and
- The generated image, audio or video content constitutes "deep fake" content (also defined in the AI Act).

Thus, to determine whether content you are generating must be specially identified, you need to resolve three key issues:
- Whether you meet the conditions to be regarded as a deployer
- Whether you are using a tool to create content which qualifies as an AI system within the meaning of the act, and
- Whether the content may be regarded as a deep fake.

Here the definitions from the AI Act will come in handy.

Under Art. 3(4), a **deployer** is "a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity."

In turn, under Art. 3(1), an **AI system** is "a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions,

content, recommendations, or decisions that can influence physical or virtual environments." Thus, this does not mean just any IT system, but only one that meets the conditions set forth in the definition.

And under Art. 3(60), for content to be regarded as a **deep fake**, it must be "[i] AI-generated or manipulated [ii] image, audio or video content that [iii] resembles [iv] existing persons, objects, places, entities or events and [v] would falsely appear to a person to be authentic or truthful."

A major element of Art. 50(4) of the AI Act is the "**artistic clause**," which somewhat softens the requirements for identifying artistic content, but without entirely excluding the labelling obligation. The clause reads: "Where the content forms part of an evidently artistic, creative, satirical, fictional or analogous work or programme, the transparency obligations set out in this paragraph are limited to disclosure of the existence of such generated or manipulated content in an appropriate manner that does not hamper the display or enjoyment of the work."

More detailed clarifications on the identification of content are set forth in the Guidelines on the implementation of the transparency obligations for certain AI systems under Art. 50 of the AI Act, published by the European Commission on 20 July 2026 (the "**Transparency Guidelines**").

Another document worth consulting in this area is the Code of Practice on Transparency of AI-Generated Content (the "**Code of Practice**"), drawn up by EU working groups, which contains detailed guidance on technical aspects of labelling of content. The Code of Practice is not binding law, and there is not obligation to comply with it. Still, it is worth considering formally signing onto the Code of Practice, as following it provides a degree of certainty that your actions comply with Art. 50(4) of the AI Act.

To bring these introductory remarks to a close, we should mention that the failure to properly designate materials generated through an AI system can be subject to a very high potential fine, of up to EUR 15,000,000 or 3% of an undertaking's worldwide annual turnover for the preceding financial year, whichever is higher (AI Act Art. 99(4)). This is another reason to take the transparency obligation very seriously.

The transparency obligation enters into force on 2 August 2026.

Q 1: Do the transparency duties provided in the AI Act apply to materials created outside the EU or distributed globally, if they also reach recipients in the EU?

Yes. In the AI Act, what is basically decisive is not the place where content is created, but the location of the **intended result (output)**. Thus the AI Act applies to materials created outside the EU and to AI systems located in a third country, if the results (output) are "intended to be used in the Union."

The phrase "intended to be used in the Union" allows a certain margin for interpretation. But if content is actively directed toward the EU market or it is reasonably foreseeable that the content will reach EU recipients, then the AI Act applies, including the transparency duties. It follows that materials distributed globally will also be covered by the AI Act insofar as they reach recipients in the EU.

It should be remembered that the aim of the AI Act is to provide effective protection of recipients in the EU, which requires combating attempts to circumvent the regulations. Thus the AI Act also applies to entities from third countries and the content they create, if it is intended to reach the EU. This concept intentionally tracks the same logic used under the General Data Protection Regulation, which is also widely known.

Q 2: What about materials published before 2 August 2026, but broadcast, extended or reused after that data?

The Transparency Guidelines offer advice on interpreting this issue, stating that materials generated and published prior to 2 August 2026 are not formally subject to the transparency obligation. Nonetheless, the guidelines encourage labelling of such materials by the entities that generated them or have them in their possession.

However, if materials were generated before the effective date but not published before then (i.e. the initial publication occurs on or after 2 August 2026), then the content should be labelled in compliance with the AI Act.

Q 3: How to treat hybrid materials, for example using real actors or models plus an AI-generated background, staging or individual elements of AI?

To answer this question, it must first be determined whether it involves a deep fake.

A deep fake is image, sound or video content, generated or manipulated by AI, that resembles existing persons, objects, places, entities or events, which a recipient could incorrectly regard as authentic or truthful.

Four conditions must **all be fulfilled** for an AI-generated image, sound or video to be regarded as a deep fake:

- The image, sound or video "resembles"—i.e. is similar to but need not be identical to…
- "Existing"—understood to mean that the content of the image, sound or video resembles people or things that exist or could exist in reality…
- "Persons, objects, places, entities or events"—here the notion of "persons" is understood to mean realistic humans (including digital replicas of real people, realistic avatars or human characters generated by AI, as well as characteristic features such as appearance, voice, behaviour etc)…
- Which could "falsely appear to a person to be authentic or truthful"—understood to mean misleading the audience as to the truth of the content, image or video.

Deep fake	Not a deep fake
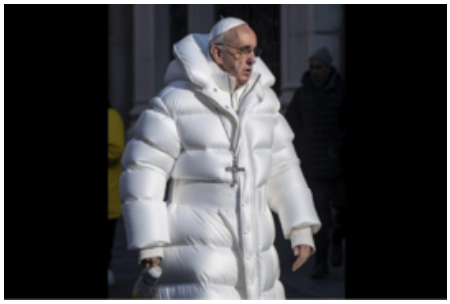	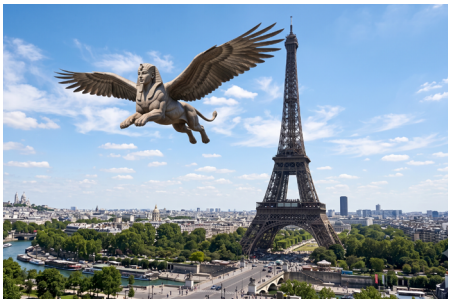
Image generated by Midjourney (source: Forbes). It resembles an existing person (the late Pope Francis) but gives the audience a false impression of authenticity—they might assume that in the winter, the pope actually wore a papal-length puffer jacket.	Image generated by Gemini (our own). It resembles an existing place (Paris), but does not create an impression of authenticity for audiences, because sphinxes do not exist, not to mention sphinxes flying over Paris.

Deep fake	Not a deep fake
An AI-generated video depicting a person resembling a specific politician giving a speech before an audience.	A cartoon generated by AI based on an existing image depicting a historical event.
An AI-generated video depicting an AI-generated influencer or celebrity in an advertising or promotional context.	An AI-generated video depicting mice arguing in human language over which cheese is the best, as part of an ad campaign by a cheese producer.

For hybrid materials, i.e. combining AI-generated elements with real objects or real (live) models, what is of primary importance is the possibility of creating a false impression of authenticity.

Unfortunately, there is no agreed measure or proportion determining "how much AI" in an image or video makes the item a deep fake—each case must be assessed individually.

Below we discuss a few instances of hybrid materials:

Real person + AI-generated background or setting

The mere fact that a person (the model) is real doesn't exclude the labelling obligation for deep fakes. What is relevant is whether the viewer of the image or video could be misled as to the authenticity of the scene depicted.

For example: a model + a realistic (AI-modified) backdrop of a café in Paris would probably constitute a deep fake requiring labelling, whereas a model + a futuristic (AI-generated) sci-fi background (clearly fictional) probably would not.

A few minor AI elements

If the AI-generated element is marginal, doesn't alter the message, and doesn't affect the perception of authenticity, the argument for applying the labelling obligation under the AI Act is weaker, but not ruled out completely. It would be different if a specific element were crucial for the message—for example, an AI-generated weapon held by a person, a crowd of demonstrators, wartime destruction, or exotic location. In such cases, even "just one element" could totally change how the content is perceived and give rise to the obligation to label the content appropriately.

Q 4: When does the use of AI in post-production constitute content generated or modified by AI, and when is it just standard processing?

The AI Act is highly enigmatic in this respect.

The AI Act does not say which activities involving editing of an image (e.g. increasing the resolution, sharpening the image, erasing elements in the background, or correcting the colours) makes the image, sound or video content into a deep fake, although the definition of "deep fake" does provide some clues. For a deep fake to arise, the editing must result in generation or manipulation of the content. Manipulation of content would consist of falsifying the original content in a manner suggesting a situation different from the original one, or altering the meaning or emotional charge for the audience.

- Simple editing limited for example to removing the "noise" from an image, to a degree that doesn't add new elements to the image which did not potentially exist in the original image, **does not create a deep fake**.
- Removing people from a photo (leaving aside the artistic clause for the moment), insofar as it does not result in falsifying the message flowing from the photo, **does not necessarily create a deep fake**. But, for example, erasing people from a beach so that there is no one in the background **could constitute a deep fake**, if it creates the impression that the site is empty when in reality it is usually crowded.
- Correcting colours will not constitute a deep fake so long as it does not alter how the content is received (e.g. "punching up" the saturation of the landscape to enhance its beauty, or using an LUT for colour grading to match the rest of a film). An example of a deep fake in this respect might be in a photo or video of a meeting of politicians, where the colours are adjusted to create a sense of danger or a hostile atmosphere, when the real situation was clearly different, or enhancing a politician's features to give a false impression of his or her state of health.
- A problematic issue is using AI to "expand" images. It would be safe to assume that expanding the setting (when it is not easily recognisable, for example meadows) is permissible so long as it involves adding similar elements (more meadows, more woods, more of the identical texture in the background). This sort of editing does not constitute a deep fake. However, it might constitute a deep fake to expand an image to include new elements that could mislead the viewer (e.g. expanding a view of meadows to make a landmark like the Wawel Castle visible in the image, when it wasn't there before).

Evidently, this is not a black-and-white situation. The context of the actions and the degree of the resulting changes must be taken into account.

Q 5: How should AI content be identified depending on the format: videos, internet, TV, radio, a song, a music video?

Under the Transparency Guidelines, regardless of the format used, the labelling should be legible and clear (whether audio or visual) and ensure that the recipient can learn that the content is a deep fake from the moment of first contact with the content. The information that a fragment constitutes a deep fake should be conveyed in a way that stands out from the other information and the environment in which the content is presented.

It must be remembered that content may reach audiences of various ages, and thus you need to ensure that the labelling is legible and clearly understandable for groups such as children and the elderly.

Under the Transparency Guidelines, labelling of a deep fake will not be regarded as clear if the audience could easily miss or ignore the information during normal display or interaction with the content. An example given of inadequate labelling would be identifying the content as a deep fake only in the terms and conditions for use of a website.

Examples of icons that can be used to identify content generated or modified by AI

With respect to content to which the artistic clause does not apply, the labelling could be presented as follows.

Video content should be tagged from the first frame of the deep fake, ideally for its entire duration because the viewer might scroll through the material (if the form allows) or enter the space while it is ongoing. You should avoid situations where the viewer's first contact with the material happens where there is no information displayed stating that the video was created or manipulated by AI.

In this respect the Code of Practice recommends placement on the screen of an icon with the acronym "AI" (a standard form or your own, if it is sufficiently legible). Depending on the situation, it is recommended to add "generated" or "modified" (in English or in the language in which the content is displayed). There are no required dimensions for the icon, but it must be legible.

On platforms with interactive options, it is recommended to add a "second layer" where the user can, for example, hover over the AI acronym with the mouse to expand the acronym to explain which moments or scenes were modified and in what respect why were modified.

In the case of live transmissions containing deep-fake elements, the labelling should appear on the screen, at the latest, at the moment when the deep fake appears. It is also worth adding information (audio and visual) at the start of the broadcast warning that certain passages will contain deep-fake content.

Images—placement analogous to video content.

Audio—a short audible disclaimer should be used before the start of content constituting a deep fake. The message should be in the same language as the deep-fake audio content. In the case of longer content, inserts may be made during the deep fake. It is also worth adding a disclaimer after the end of the piece.

If the deep-fake audio content is published in a medium also allowing the transmission of images, the Code of Practice recommends adding a visual indicator. Remember that the label must be legible. For example, an "AI" icon on the cover of an AI-generated audio album might suggest that only the cover was generated by AI. If the platform allows, an indication may be added in the name or using the mechanisms provided by the specific platform (without additional labelling).

The European Commission has proposed icons for labelling content—a complete set can be found at EU Icons for labelling AI-generated content | Shaping Europe's digital future.

Q 6: Can the entire material be labelled as AI if only some of the shots or elements were generated or modified by AI?

As a rule, it is the content that constitutes a deep fake that should be labelled as such. But in the case of brief or dynamic content, such labelling might not be legible for the audience. Thus, if not misleading to the audience, the entire piece or longer sections may be tagged.

If it is feasible to use a separate layer (interactive, integrated with the playback), detailed labels should be used indicating which scenes and elements were generated or modified by AI, or vice versa—indicating the scenes in which AI was not used, if they are in the minority. When the use of a separate layer is not feasible, a table could be placed at the end of the piece explaining this issue.

Q 7: Who is responsible for labelling and legal compliance: customer, agency, studio, producer, influencer or publisher?

The entity responsible for labelling AI content is the deployer, i.e. the entity using an AI system under its authority to generate content. Under the usual model, where the customer hires the agency to prepare advertising materials, and the agency carries out the assignment using its own employees or associates, the deployer will typically be the agency. It is the agency that coordinates the preparation of the material and is responsible to the customer for preparing the material. The agency also takes the decision to use an AI system and oversees its use.

It is much harder to determine whether, in this scheme, the role of deployer may also be assigned to the agency's employees/associates and the customer. The Transparency Guidelines are favourable to agency staffers in this respect, stating that where the deployer is a legal person under whose authority the AI system is used, such as an advertising company, the individual employees actually using the AI system for this purpose, but acting under the deployer's instructions and control, are not considered separate deployers. Unfortunately the guidelines don't expressly address the case where creators are affiliated with an agency under a non-employment contract. In this case there are a range of analogies with the case of an employee, but the risk that such creators would be regarded as deployers is higher.

And at first glance, the customer is not using the AI system or exercising control over it. From this perspective, the customer should not be regarded as a deployer. The Transparency Guidelines lean in this direction, stating that "a

company that *merely commissions* an advertising agency to produce an advertisement, without taking decisions and exercising control over whether and how the advertising agency uses AI in the production process, is not a deployer." Nonetheless, we suggest caution in this case. In our view, it cannot be ruled out that in certain instances the customer could also be regarded as a deployer, particularly if it is directly involved in creating the material, including taking day-to-day decisions on the shape of the material and the tools used to create it.

The deployer must label the material appropriately, and it is also the deployer against whom an administrative fine for improper labelling may be assessed. Material not compliant with the labelling requirements automatically becomes unlawful material, which could have a range of other implications (for example, the material could be blocked by online platforms, or improperly labelled advertising could be regarded as an act of unfair competition).

In addition, deployers who participate in complex chains for production and distribution of content should take appropriate measures to ensure that the labelling of content required under AI Act Art. 50(4) is clearly visible and easily recognisable by audiences when they first come into contact with the content. In practice this could additionally require, for example, the inclusion of appropriate provisions in contracts with distribution partners.

Q 8: Who will oversee the use of AI, how will it be possible to verify how material is created, and who will impose fines?

In Poland it will be the Commission for Development and Safety of Artificial Intelligence, a new administrative body to be established for supervision of AI systems (as provided in the AI Systems Act of 3 July 2026).

Q 9: Can avatars or AI-generated characters recommend products which they could not have tested in reality?

The AI Act does not prohibit avatars or AI-generated characters from recommending products, but only imposes a duty of transparency. The issue of the truthfulness of the recommendations as such is governed separately by consumer protection law.

When is an AI avatar a deep fake requiring labelling?

The AI Act's definition of a deep fake covers not only digital replicas of specific, existing people (e.g. celebrities), but also realistically rendered human avatars, even depicting people who don't actually exist. What is decisive is whether the recipient could falsely regard the person as real.

It will almost certainly be regarded as a deep fake if the avatar:

- Looks like a real person (with a realistic face, voice, body), and
- Appears in a context of advertising or promotion, e.g. speaks to the camera about the advantages of a product, reviews the product, or encourages users to buy the product.

The Transparency Guidelines give as an example of a deep fake, covered by the transparency obligation, a video with an AI-generated image of a celebrity influencer in an advertising or promotional context. Then the material must be clearly and legibly tagged, at the latest, when the user first encounters the image.

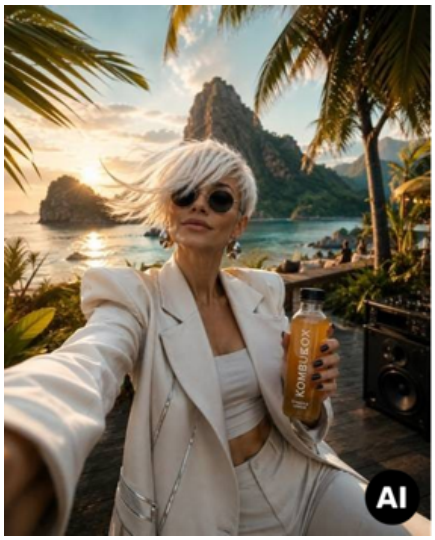

AI-generated influencer advertising kombucha (our own project)

The AI Act and consumer protection law—two independent regimes

The AI Act does not regulate whether what the avatar says about a product is truthful. This issue is covered by other laws, including consumer protection law.

What is primarily problematic in this respect is the suggestion of personal experience which an AI could not have had (e.g. "I tried this cream," or "These are my favourite earphones"). Such statements made by a fictional character could be misleading to audiences. An AI influencer could recommend products based on data, but should not present non-existent personal experiences as facts.

Q 10: How does the artistic clause impact the AI transparency obligation?

If a deep fake is part of a work of a clearly artistic, creative, satirical, fictional or analogous nature, the transparency obligation is softened but not eliminated entirely.

According to the Transparency Guidelines:

- **Artistic works** are works created for the purpose of art (music, film, visual arts)
- **Creative works** are works that involve creative choices (this excludes works mainly motivated by functional or technical considerations)
- **Satirical works** are works intended to criticise society, politics, business or public figures using humoristic techniques (irony, sarcasm, pastiche etc)
- **Fictional works** are cultural works placed in an imaginary setting
- **Analogous works** are works sharing core traits with the above categories, but not fitting neatly into one category.

Content could qualify for more than one of these categories, and thus a case-by-case analysis is required.

Significantly, for the artistic clause to be applicable, it must be "evident" to the audience that the work falls into one of these categories.

An example of an artistic work would include AI-generated special effects in film scenes constituting deep fakes, such as:

- "De-aging" existing actors
- Digital replicas of dead people

- An AI-modified image of an existing politician in a scene clearly aimed at criticising political decisions taken or supported by the person.

However, AI-generated images of well-known people suggesting their involvement in certain activities, but lacking any fictional, satirical or analogous character, will not qualify for any of the aforementioned categories.

If it would not be clear to the audience that the main aim of a work is to convey content that is artistic, satirical etc, the work is not protected by the artistic clause.

But if there are doubts whether the artistic clause is applicable to a work, then, instead of using a full, visible label, it is sufficient to disclose the **existence** of AI-generated or -manipulated content in the appropriate manner, without **interfering with the display or enjoyment of the work**.

The place and manner of disclosing information about AI content depends on the medium and format of the work. The only condition is that the labelling is noticeable to the audience but does not disrupt the experience of the work itself. The assessment of whether the manner of disclosure is appropriate in the given case must be determined in each instance by the nature of the work, the distribution channel, and the audiences' expectations.

In every case, the audience must have a realistic opportunity to note this information, before or at the time when they encounter the work. For example:

- In a film at the cinema, in the end credits or accompanying materials accessible to audiences (e.g. the cinema programme or the distributor's website)
- In an exhibition at an art gallery, on a notice board at the entrance, or next to the item itself, if clearly visible to visitors
- In a video game, in the product description at the point of sale, in the credits section, or on an informational screen displayed at startup.

Crucially for the advertising industry, the Transparency Guidelines **exclude reliance on the artistic clause** for content solely of an informational or **commercial** nature. Consequently, the protection of this clause will not be available for most advertising.

Authors

Julia Padlewska-Oleksy
attorney-at-law
intellectual property

Monika A. Górska
attorney-at-law, partner
intellectual property

Lena Marcinoska-Boulangé
adwokat, partner
intellectual property

Bartosz Kurzec
younger associate
new technologies law

Krzysztof Wojdyło
adwokat, partner
new technologies law

Marysia Surzycka
CEO & Head of Production
@surzycka

Szymon Furi Brzozowski
AI Lead
@furi.user

Maciej Miliszkiewicz
AI Artist & educator
@mmiliszkiewicz

WARDYŃSKI & PARTNERS

SHOOTME