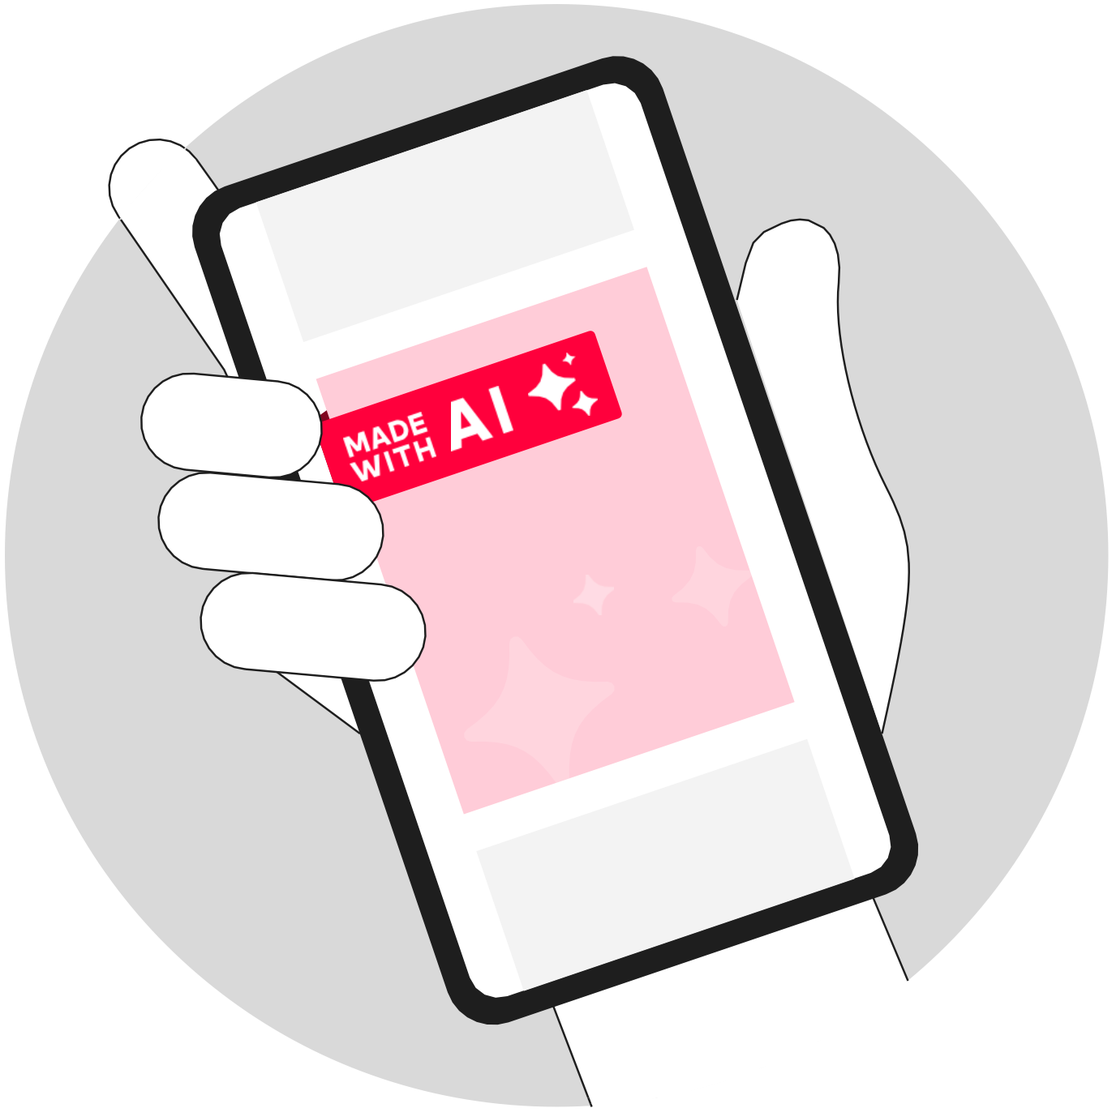
The EU AI Act: a practical guide for brand marketers

Navigating the AI labelling obligations in practice



In collaboration with:

LEWIS SILKIN

September 2026

Contents

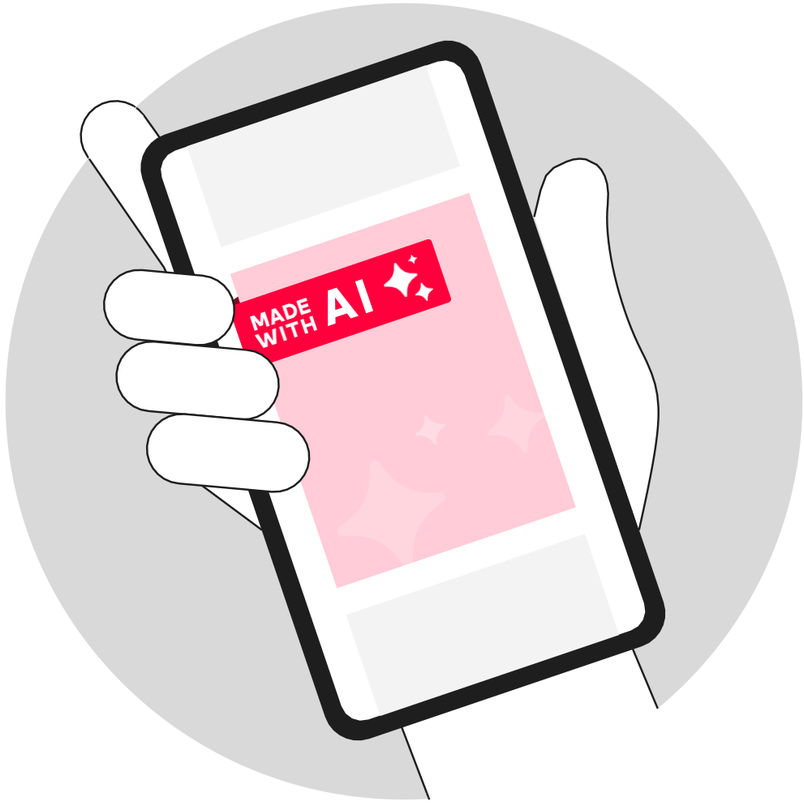
Introduction 3

Breaking down the EU AI Act 4

- What does the Article 50(4) labelling obligation require? 4
- Who is caught by this obligation? 4
- Who is a deployer in the advertising supply chain? 5
- What is a deep fake? 6
- Can you rely on an exception for creative content to avoid the disclosure obligation? 7
- Do I need to retroactively label content created before 2 August 2026? 7
- Do you have to label AI-generated text in advertising copy? 7
- If your content meets the “deep fake” test, how should you label it? 8

What brands should do: Governance and contractual action points 9

Appendix: Practical scenarios 10



Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

Introduction

78% of brands are now using AI-generated and AI-enhanced content in their user-facing marketing creative. But as AI becomes increasingly embedded in advertising, marketers are facing new questions about **when and how AI use should be disclosed, who is responsible, and what good governance looks like.**

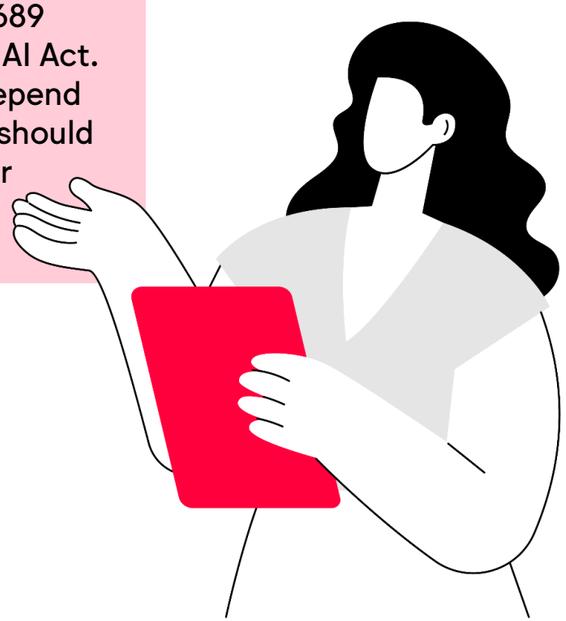
The EU AI Act brings these questions into sharper focus. **As of 2 August 2026, Article 50(4), any AI-generated and manipulated content that qualifies as a 'deep fake' will need to be labelled.** For brands, however, a lack of clarity remains around what constitutes a 'deep fake', who is responsible for disclosure and what an appropriate label looks like.

This uncertainty is reflected in a clear action gap: In July 2026, **only 25% of brands had a clear transparency policy outlining how and when their marketing teams should disclose AI use.**

With potentially significant legal and reputational consequences, including maximum fines of €15 million or 3% of global annual turnover, whichever is higher, AI labelling is becoming an important governance priority.

This document aims to help brands navigate that gap. Developed by WFA in partnership with Lewis Silkin, and building on earlier WFA global guidance on AI transparency and labelling, it summarises the requirements of Article 50(4) and offers practical considerations for marketers to take into account.

DISCLAIMER: This guide is for general informational purposes only and does not constitute legal advice. It reflects general commentary based on the Code of Practice on Transparency of AI-Generated Content, the European Commission's Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689 published on 20 July 2026 and the EU AI Act. The application of Article 50(4) will depend on specific factual scenarios. Brands should seek independent legal advice on their individual compliance obligations.



1 What does the Article 50(4) labelling obligation require?

Article 50(4) of the EU AI Act ("**AI Act**") requires "deployers" of AI systems to clearly disclose that content has been artificially generated or manipulated in two distinct scenarios:

- where image, audio or video content qualifies as a "deep fake" (explored further on page 6) which is likely to capture a variety of AI-generated advertising and creative visual assets; or
- where AI generated or manipulated text is published to inform the public on matters of public interest and has not undergone human review or editorial control.

Routine AI-drafted advertising copy will generally not require a label, as it will almost always have undergone human review and editorial control. Accordingly, this guidance note focuses on the first scenario: "deep fakes".

2 Who is caught by this obligation?

Even though the AI Act is a piece of European legislation, this doesn't mean that it only applies to organisations established or based in the EU. The AI Act has "extra-territorial" reach, meaning the AI labelling rules apply to both:

- deployers established or based within the EU; and
- deployers in third countries (e.g. outside the EU) where they foresee, direct or authorise distribution of AI content into the EU, including by posting deep fakes on the "globally accessible internet".

The EU Commission Guidelines published on 20 July 2026 (the "**Guidelines**") suggest that non-EU deployers are not expected to be caught where their content only reaches EU audiences through channels that are "unforeseeable and outside their control". Brands running campaigns with obvious EU reach (e.g. via the publicly accessible internet) should assume the disclosure rules apply, however a campaign geo-targeted to a non-EU country that reaches the EU only incidentally (e.g. shared privately via WhatsApp) is unlikely to be caught.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

3 Who is a deployer in the advertising supply chain?

"Deployers" are those who use AI systems under their "authority" and who control how the AI system is used. "Authority" is understood broadly: it means assuming responsibility for deploying the system and for the manner of its actual use (including its outputs), but this does not mean hands-on technical control over the AI system itself. It is enough that a party decides for what purposes and how the AI system is used, even across complex supply chains.

The Guidelines make clear that where the deployer is a legal person such as "an advertising company", the corporate entity will bear compliance - individual employees (digital animators, web designers, content creators) acting under its instructions are not themselves separate deployers.

For brands, even if you do not use AI tools yourself to create campaign materials, it is perfectly possible for you to still be deemed a "deployer" of AI used in your supply chain - particularly where you exercise some level of input or control over the AI tools used by your third-party suppliers and agencies. For example, you may set parameters around an agency's AI use in contractual terms or via an AI addendum. You will also typically have final approval rights over campaign content and the decision to go live with that content which includes the AI asset - making it more likely you will be considered a "deployer" under the AI Act.

Theoretically, there may be limited scenarios where a brand would not be considered a "deployer" if your role "is limited to disseminating or transmitting AI-generated or manipulated content created by third parties" - which the Guidelines expressly state would not afford "deployer" status. However, a brand's typical role in commissioning and approving AI generated content under its direction goes beyond a passive role of "disseminating or transmitting AI-generated content created by third parties", so brands should safely assume they will usually be considered a "deployer" in respect of AI tools used within their supply chain.

Your agency is also likely to be a "deployer" if it retains control and discretion over which AI tools are used and how they are applied during the creative process. The Guidelines include an advertising-specific example that reinforces this: "a company established in a third country that uses an AI system to generate a deep fake of a celebrity featured in an advertisement displayed in the Union is also a deployer falling within the scope of the AI Act". The example makes no reference to whether the company was acting on a client's instructions – the decisive factor is that it "uses" the system.

In many circumstances, both brand and agency will independently qualify as "deployers" of AI systems used to generate campaign content - meaning regulatory compliance is a shared responsibility that cannot be fully outsourced contractually. Even where deployer status is arguable, the EU Commission expects all supply chain actors to implement "proportionate and technically feasible measures to enable natural persons to detect the AI-generated or manipulated content", so brands should maintain appropriate oversight of AI-generated content in their campaigns.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

4 What is a deep fake?

Even if you are a "deployer", you are only required to label AI-generated or manipulated (visual) content that constitutes a "deep fake" (as defined in the AI Act).

"Deep fake" means AI generated or manipulated image, audio or video content which:

(1) resembles existing persons, objects, places, entities or events; AND
(2) falsely appears to a person to be "authentic or truthful".

This is a two-stage test and both limbs must be satisfied. The Guidelines confirm it is sufficient for content to resemble someone or something that "exists, can plausibly exist or could have plausibly existed in reality". Content that defies the laws of nature (e.g. pigs flying) falls outside scope. Because this first limb of the test is presented so broadly, it will be easily satisfied so in practice this means the assessment of whether something constitutes a "deep fake" requiring a label will often turn on the second limb of the test: does the content falsely appear to a person to be authentic or truthful?

The "falsely appears" assessment is objective and does not require intent to deceive. Context will be crucial when it comes to assessing whether the "falsely appear" test is met. It must be assessed as a whole, considering: the level of resemblance to the real subject, the substantive message, the intended and foreseeable deployment context, the environment of presentation, and the intended and reasonably foreseeable audience. If the reasonably foreseeable audience is likely to consist of vulnerable groups (e.g. children, the elderly or those with accessibility needs), this will lower the threshold at which content qualifies as a deep fake.

In an advertising context, the key determining question for assessing if the second limb of the test is met is: **does this AI-generated or manipulated element affect the audience's perception of the authenticity and truthfulness of the advertisement and product?**

The Guidelines state that "AI-powered colour correction, background extensions of existing content, adjustments or replacements of backgrounds for clearly aesthetic purposes, compositions and arrangements of existing products, or re-scaling of images applied in product advertisements or packaging is likely to have only a minor impact on a person's perception of the authenticity and truthfulness of the advertisement and the product" - meaning it would likely fail the second limb of the "deep fake" test.

Conversely, an AI-generated image of a product in an advertisement that makes the product appear not identical to the real product, more appealing or with improved quality than in real life would affect the audience's perception and mislead as to the actual product appearance, characteristics or use.

Put simply, an AI-generated backdrop used for cosmetic and technical purposes to set a scene is unlikely to require labelling, provided it does not affect how the audience perceives the authenticity of the actual product, person or event that is the focus of the advert. The position differs where AI use touches the item being sold or if it is a prominent feature in the context of advertisement. See the **Practical Scenarios in the Appendix** for more information on how the deep fake test might be applied in practice.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

5 Can you rely on an exception for creative content to avoid the disclosure obligation?

The AI Act provides a lighter transparency regime for content that is "evidently" artistic, creative, satirical or fictional. However, the Guidelines are clear that content whose nature is "exclusively informative or commercial and is recognisable as such" is excluded from this lighter regime. The commercial nature of advertising means that any advertising content including deep fakes will usually face the full labelling obligations.

6 Do I need to retroactively label content created before 2 August 2026?

AI deep fakes created before 2 August 2026 do not need retroactive labelling. However, the EU Commission makes it clear that labelling would be deemed "good practice" if the deep fake is published after 2 August 2026, particularly for clear-cut deep fakes where labelling will not require "disproportionate efforts" (which is unlikely for digital assets).

7 Do you have to label AI-generated text in advertising copy?

AI-manipulated text in advertisements does not need to be labelled **unless** it includes claims related to matters of public interest e.g. health, consumer safety or sustainability. Brands in sectors such as food, pharmaceuticals, cosmetics, financial services or energy should pay particular attention to this, as marketing copy in these industries frequently touches on public interest matters.

However, if the AI-generated text in question has undergone human review or editorial control (even if it does relate to public interest matters), it does not need to be labelled (provided there is no further substantive AI intervention afterwards). Human review means examining content substance and exercising professional judgment; it excludes superficial checks (i.e. spell-check). Editorial control refers to a person holding ultimate legal responsibility over the substance of the content and publication (e.g. an individual, editorial board, or publishing company). In practice, advertising copy will almost always be subject to human review, so the text labelling obligation will rarely be triggered for standard advertising.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

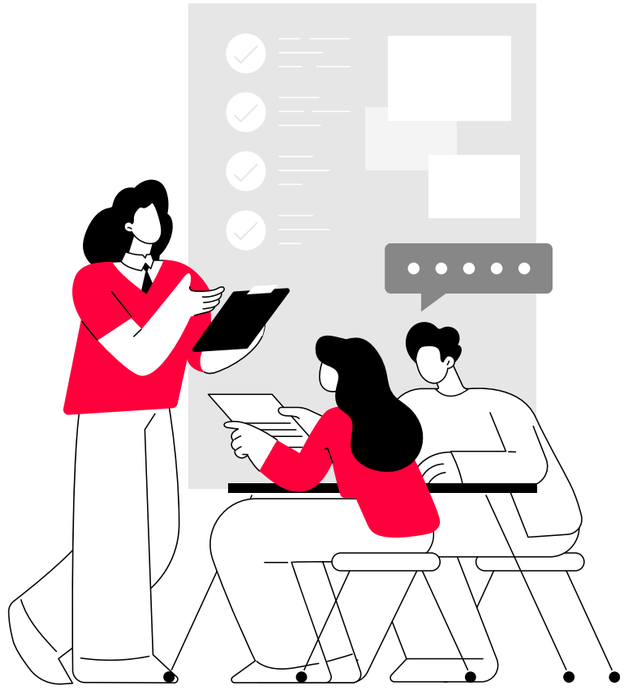
8 If your content meets the “deep fake” test, how should you label it?

The disclosure must be **clear, distinguishable and understandable to the audience**, factoring in whether children, the elderly or those with accessibility needs might be exposed. It must appear when the content **first reaches its audience**, not buried in small print or hidden in metadata.

There is no single mandated format. The European Commission has recommended an EU icon (a simple “AI” mark, with “generated” or “modified” as appropriate) in its Code of Practice on Article 50, which deployers may use to help demonstrate compliance. However, adherence to the Code is voluntary and is not (in of itself) conclusive evidence of compliance, and use of the EU icon is not mandatory.

You may have preferences on AI labelling (and placement) in your campaign assets, but any label must be noticeable and clear while preserving creative quality. Most online platforms (like Instagram and TikTok) have built-in tools to label AI content. However, if the platform label is invisible (i.e. a machine-readable marking), you cannot rely on it to discharge your disclosure obligations - a human-perceivable and prominent label is still required. Even if the platform-applied label is prominent, it remains the brand’s responsibility to consider if AI disclosure is required and if the label complies with AI Act requirements.

IMPORTANT: Labelling a deep fake does not make the content lawful or exempt it from other rules. The transparency obligations sit alongside data protection law (where a real person is depicted), intellectual property and personality rights, and advertising regulation and consumer protection law. Compliant labelling addresses the AI Act’s disclosure requirement but not the underlying lawfulness of the advertisement so disclosure obligations should always be considered alongside applicable national laws.



Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

What brands should do: Governance and contractual action points

The following governance actions should be immediate priorities for brands:

Develop a clear AI policy.

Define which AI use cases are permitted across marketing operations, who has authority to approve AI-generated assets (and decide if a label is required), and the escalation path for borderline cases.

Build AI disclosure into creative approval workflow.

Agree early in the creative process (ideally at briefing stage) if a label is required, what label to apply and where. The label must be present when content is published or broadcast, not added after the campaign has gone live. This requires a systematic review point in your sign-off process: for each asset, has AI been used, does it meet the deep fake test, and is an appropriate label in place?

Include AI labelling obligations in agency contracts.

Require creative agencies to identify AI-generated content early in the creative process and require them to cooperate with you to agree and implement appropriate disclosure. Ensure agencies and distribution partners are contractually required to preserve labels through to audience exposure (if they are responsible for distributing the assets to the public). The Guidelines require deployers in complex value chains to take “proportionate measures” (including contractual conditions) to ensure labelling is preserved downstream - your contracts will help to demonstrate this compliance.

Keep records.

Create internal processes recording: (i) the decision to use AI; (ii) what content was generated or manipulated using AI; and (iii) the assessment of whether disclosure is required under Article 50(4). Consider requiring agencies to maintain a log of AI tools used and content generated. This documentation serves as both a compliance record and evidence of appropriate oversight.

Consult legal.

If you are unsure if the disclosure obligation is triggered (as each example is going to turn on the specific facts). Certain assets may involve interpretative complexity when assessed in context against the Guidelines (particularly those cases that are borderline or if they reveal “gaps” in the Guidelines).

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

Appendix: Practical scenarios

SCENARIO	ANALYSIS
✦ **AI-generated background crowd/people** A brand uses Gen AI to create realistic humans in the background of an ad. The individuals do not interact with the main product of the ad and do not influence the product/service proposition or how the product is perceived (e.g. walkers-by to make a street appear more busy).	**Likely NOT a deep fake requiring disclosure**. The AI-generated humans resemble persons that can plausibly exist (even if they are not identifiable people) so the first limb of the deep fake test is satisfied. However, if the AI-generated humans are used for purely scene-setting and cosmetic purposes and if they do not interact with the main product of the ad or influence how the product is perceived, they are unlikely to affect a person's perception of the authenticity and truthfulness of the advert and the product (so the second limb of the deep fake test is unlikely to be satisfied).
✦ **AI-generated person interacting with a product** A brand creates a photorealistic AI-generated person who does not correspond to any real individual and uses them in an advertisement to interact directly with the product e.g. eating a food product or using an electronic device or an AI-generated model trying on a garment.	**Likely IS a deep fake requiring disclosure**. The AI-generated person resembles a person that can plausibly exist (even if it does not correspond to a real individual) so the first limb of the deep fake test is satisfied. If the AI-generated person is the core endorsement mechanism and their product interaction goes to the substance of the advertising claim, the second limb of the deep fake test is likely to be satisfied. This is analogous to the "realistic synthetic influencer testing out a sponsored real product" example in the Guidelines, which attracts the full labelling obligation.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

Appendix: Practical scenarios

SCENARIO	ANALYSIS
✦ **AI-generated product variations** A fashion retailer uses Gen AI to create photorealistic versions of a garment in different colours or materials that accurately correspond to variants genuinely available for purchase, without photographing each physical variation.	**Likely NOT a deep fake requiring disclosure**, provided the generated product variations accurately represent the real product variants available for purchase. If the AI content does not mislead a person as to the product's actual characteristics, disclosure is unlikely to be required. However, if AI-generated variations depict the garments as looking not identical to, or materially better than, the real article, the analysis would change and disclosure may be required. The key question in this scenario is whether the AI-generated variations could mislead the audience about the garment and its characteristics. This will need to be assessed in context of the advert, depending on the facts.
✦ **AI-generated recreation of an existing place or event** A brand uses Gen AI to create a photorealistic image or video of an existing place (e.g. a Parisian street) or to recreate a real-world event. The place/event exists or existed, but the specific scene depicted including the people, activity or circumstances did not actually occur.	**Requires careful assessment**. If the place/event existed or could plausibly exist, then this would satisfy the first limb of the deep fake test but the key question is whether the fabricated scene could mislead the audience about the "nature of the advertisement and the product" (to understand if the second limb of the deep fake test is satisfied). If the AI-generated place is merely a cosmetic adjustment or enhancement, disclosure is less likely to be required. If the place IS the product (e.g. destination marketing) or if it is material to the advertising claim, disclosure is more likely required. This will need to be assessed in context of the advert, depending on the facts. For image and video content, the Guidelines acknowledge that a high degree of photorealism renders it more likely that content will be considered a deep fake (as it likely resembles existing subjects), but photorealism alone is not determinative for the assessment. What should still be assessed, is whether such content is capable of deceiving or misleading a person regarding the content's authenticity or truthfulness.

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.

About WFA and Lewis Silkin

About WFA

WFA champions more effective, efficient and responsible marketing communications which help WFA members drive brand growth. It is the voice of marketers worldwide, representing more than 150 of the world's biggest brand owners and national advertiser associations in over 60 markets, which speak for tens of thousands of brands at a local level. Through WFA, they form a global network that offers marketers a unique source of leadership, expertise and inspiration.

For more information, please contact:



Gabrielle Robitaille
Director of Policy
g.robitaille@wfanet.org



Maxime Ricard
Senior Digital Policy Manager
m.ricard@wfanet.org

About Lewis Silkin

Widely recognised as the UK's leading advertising and marketing law firm – with more ranked experts in this field than any other firm and rated at the highest level by both Chambers and Legal 500 – Lewis Silkin's advice has shaped legal precedents and helped transform industries.

With unrivalled depth across all aspects of traditional and digital marketing, regulatory compliance, brand protection, and the full range of advertising sector agreements, Lewis Silkin advises brands, agencies and platforms on the issues that matter most, from AI and ad tech to content labelling and consumer regulation. As the only UK law firm appointed to GALA (the Global Advertising Lawyers Alliance), Lewis Silkin offers international reach across more than 90 jurisdictions with a team of more than 950 people working across nine offices.

For more information, please contact:



JJ Shaw
Partner
jj.shaw@lewissilkin.com



Roch Glowacki
Partner
roch.glowacki@lewissilkin.com



Gemma Jones
Managing Associate
gemma.jones@lewissilkin.com

Please note nothing in this document should be taken as legal advice. This document represents one interpretation of the EU AI Act Article 50(4). Brands should seek their own legal guidance for compliance support.



World Federation of Advertisers
London, Brussels, Singapore, New York

wfanet.org

info@wfanet.org

+32 2 502 57 40

linkedin.com/company/wfa
youtube.com/wfamarketers

WFA Competition law compliance policy

The purpose of the WFA is to represent the interests of advertisers and to act as a forum for legitimate contacts between members of the advertising industry. It is the policy of the WFA that it will not be used by any company, or individual, to further any anti-competitive or collusive conduct, or to engage in other activities that could violate any antitrust or competition law, regulation, rule or directives of any country or otherwise impair full and fair competition. The WFA carries out regular checks to make sure that this policy is being strictly adhered to.

As a condition of membership, members of the WFA acknowledge that their membership of the WFA is subject to competition law rules and they agree to comply fully with those laws. Members agree that they will not use the WFA directly or indirectly, (a) to reach or attempt to reach agreements or understandings with one or more of their competitors, (b) to obtain or attempt to obtain, or exchange or attempt to exchange, confidential or proprietary information regarding any other company other than in the context of a bona fide business or (c) to further any anti-competitive or collusive conduct, or (d) to engage in other activities that could violate any antitrust or competition law, regulation, rule or directives of any country or otherwise impair full and fair competition.