

AI Images & the EU AI Act

From August 2nd 2026, EU transparency rules apply. But: Not every AI image needs a label. There are exceptions for specific situations.

➤ [EU AI Act \(Regulation 2024/1689\)](#)

➤ [Article 50: Transparency Obligations](#)

Why label at all?

A business should label their AI images, so customers can make an informed decision. Beyond that, it protects your business because competitors and other entities may report you.

What happens if you don't label?

Two risks run in parallel: reputational damage once it comes out, and fines of up to €15m or 3% of worldwide turnover under [Art. 99 of the AI Act](#)

MY TAKE ON PENALTIES

The AI Act ties fines to how serious the violation is – they must be proportionate and dissuasive. A deepfake of a real person is a different league from a generated background. It also counts whether it was deliberate or careless, and whether you fixed it once told. Company size matters too, with lower caps for small businesses. My guess: for a random marketing image, the real risk is a cease-and-desist letter from a competitor, not a big fine. *That is my interpretation, not legal advice!*

The AI Act is an addition, not a replacement

The AI Act adds one duty: disclose when content was made by AI. Existing law continues to apply in full – [copyright \(A\)](#), [consumer protection \(EU\)](#), [GDPR \(EU\)](#), [personality rights \(A\)](#) and [trademark law \(A\)](#).

Not all of this is harmonised at EU level – rules like the right to one's own image are national and vary by country.

EXAMPLE #1

Commercial use of a night-time image of the [Eiffel Tower requires permission](#), because the illumination is treated as a protected work. In [Italy, reproducing public monuments](#) for commercial purposes requires permission. Generating these images with AI instead of photographing them does not change the fact that you need permission!

EXAMPLE #2

You need an employee's consent to show them on your website or social media at all. That duty existed long before AI and comes from data protection and the right to one's own image. AI adds a second layer on top: consent to being photographed is not consent to being edited, moved into another scene, or generated. **So an image can carry a correct AI label and still be unlawful**, because the person never agreed to it.

How to label

Use the official EU icons (free, no attribution needed, [SVG](#) and [PNG](#)). Place them on the image – clearly visible, high contrast, long enough to notice. Also, make it accessible for screen-readers.

Hidden machine-readable marks ([C2PA](#), [SynthID](#)) do not replace a visible label.

A note buried in the imprint or privacy policy does not suffice.

EU AI labels

AI **AI GENERATED** **AI MODIFIED**

➤ [Download EU Labeling Icons](#)

Special case:

Art, satire, and fiction.

The rule is more flexible here: disclose that AI-generated content is present without spoiling the work. However, the line between art and advertising can be thin, and a campaign framed as creative expression may use this exception to avoid a clear label.

Decision Reference

In general: The duty of labeling applies to published work = what the public (for example your customers) sees. Internal presentations, mood boards and working drafts don't necessarily need labeling. So if you want to publish work containing AI use the following table to determine if a label is required or not.

TECHNICAL EDITS

Case	Example	Label	Why?
Basic adjustments	Contrast, saturation, white balance, sharpening on a real photo	No	Changes appearance, not content
Noise reduction	Cleaning up a low-light product shot	No	Nothing added or removed
Cropping, straightening, rotating	W	No	Framing only
Dust and spot removal	Removing specks from a dirty lens	No	Restores what was actually there
Compression	Reducing file size for web use	No	Technical only
Upscaling	Enlarging an old catalogue photo for print	No	More pixels, same content. Check that the tool has not invented new detail!
Cut-out on neutral background	Product freed from its background, placed on white	No	Standard practice, nothing invented

CLEARLY UNREAL

Case	Example	Label	Why?
Illustration or cartoon	A drawn mascot, a comic-style T-shirt print	No	Nobody will think it is a photograph
Fantasy or impossible scene	A unicorn using your product	No	Cannot be mistaken for reality
Invented, stylised place	A dreamlike landscape that clearly does not exist	No	No real reference

GENERATED CONTENT

Case	Example	Label	Why?
Outpainting / extending a scene	AI fills in the beach either side of a real photo	Usually yes	New content is generated. Small fixes a retoucher could have done by hand – removing one cloud or a mark – stay closer to editing.
Generated background, real product	Trail-running shoe placed on a mountain summit	Usually yes	A blurred or neutral backdrop makes no claim. A setting that suggests capabilities of the products does.
Season or condition swapped	Summer photo converted into a winter scene	Yes	Sells conditions that were never there.
Interior or furniture staging	Sofa shown in a photoreal AI living room	Usually yes	Looks a photograph. Real risk when the staging distorts scale or the product does not exist yet.

MISREPRESENTING THE PRODUCT

Case	Example	Label	Why?
Object made better than reality	Hotel room rendered brighter, larger, freshly renovated while the actual room is shabby	Yes	Misrepresents the exact item being sold.
Environment implying false properties	Phone shown submerged when it is not waterproof	Yes	Suggests capabilities the product lacks.
Colour altered beyond the truth	Item color shifted to a richer shade than the item shipped	Yes	No longer matches the real product.
Before-and-after claims	Diet pill ad, the „After“ image shows a slimmed down version of the person generated with AI	Yes	Fakes a real result.

PEOPLE

Case	Example	Label	Why?
Fully synthetic people	AI generated people on an ad.	Yes	Viewers assume a real person.
Real person moved into a fake setting	Studio shot of a model composited onto a mountain trail	Usually yes	An impossible scene, e.g. a person on mars, deceives nobody. A plausible one becomes a claim about the product. Consent of model is required either way.
Virtual influencer	A photorealistic AI persona recommends a product	Yes	Fakes a genuine recommendation.

Staged backgrounds and mockups

Putting a sofa into a fake room or adding a design to a coffee cup is nothing new. Designers have done it for years, and nobody feels tricked. So why is it a problem now? The practice did not change — the tool did. If an artist builds the room in 3D or a photographer turns a picture into a mockup, you don't need a label. If AI creates a room or generates a mockup, you need one. The rule looks at how the picture was made, not at whether anyone was fooled. That feels odd. But a law that had to prove harm every single time would never work, so it draws a simple line and accepts that harmless generated pictures may get flagged too.

A note on stock platforms

Buying a generated image on a stock platform also transfers the duty to label it to the buyer. If you publish a photorealistic AI image, you are the deployer and the obligation is yours even if you haven't generated the image yourself.

FAQs

Transparency obligations under Article 50 of the AI Act

[↗ Check them out](#)

What to do if you think an unlabeled image is AI?

1 Spotting it

Posters, print campaigns, YouTube pre-rolls

Small details still give it away: hands, text on signs, shadows falling the wrong way — but these tells are disappearing fast, so treat a hunch as a reason to check, not as proof. And no: anything generated before **August 2nd 2026 does not count. It is not assessed retroactively.**

2 Checking it

There are different tools you can use

↗ C2PA	who made it, what tools touched it, whether AI was involved
↗ SynthID	Google's invisible watermark in Imagen, Veo, Gemini output
↗ Metadata2Go	EXIF/IPTC data — camera, software, source type
↗ IPTC	Chrome /Firefox Browser Metadata Inspector.
↗ Google Images	Reverse search — if nothing turns up, the image likely has no history
↗ TinEye	Reverse search — if nothing turns up, the image likely has no history

3 Reporting it

This is currently the hardest part

Marketing images are not high-risk AI. They fall under the transparency rules, which means there is no incident-reporting route like the one that exists for high-risk systems. And [the supervisory authority has not been set up yet](#) — the rules apply from August 2026, but nobody is in place to enforce them. The KI-Servicestelle at RTR exists, but it is an information and contact point, not a regulator. Each country handles this differently, so check where you are.

One more thing on reporting: The AI Act applies EU-wide, so the August 2nd 2026 cut-off is the same everywhere — but the routes for complaining, and the national rules on personality rights and advertising, differ from country to country, so check where the campaign is actually running!

BEFORE YOU REPORT

Save the evidence! Download the file and take a screenshot of the media in context. Note the URL and the date. **Metadata often disappears the moment something gets re-uploaded — and an unlabeled image usually cannot be proven after it has been quietly taken down.**



About the author

Christina Bernhard MA advises on AI-generated visual content. She spent years as a designer and illustrator before moving into content strategy, and now works at the intersection of image production and technical compliance — provenance metadata, labeling standards, and the workflows that make them hold up.