

AI assistant transparency report

ARTICLE 50 OF REGULATION (EU) 2024/1689

Article 50 of the EU AI Act has been enforceable since 2 August 2026. It requires that people know when they are talking to an AI, and that the content an AI generates is marked as such. This site runs a conversational assistant, so it applies. What follows is what the law asks for, what the site does about it and which file each measure lives in, so it can be checked rather than taken on trust. I am not a lawyer and this is not a legal opinion, but my own reading of the law and what I have built from it.

What the law requires, and since when

Regulation (EU) 2024/1689 has been phasing in since 2024. The phase that reaches an assistant like this one is Article 50, on transparency obligations, and its date was 2 August 2026. The simplification package known as the Digital Omnibus, Regulation (EU) 2026/1744, pushed the high-risk obligations to December 2027 but left Article 50 untouched and on schedule. On 20 July 2026 the European Commission adopted the guidelines that interpret it and endorsed a voluntary code of practice on marking AI-generated content.

In Spain these obligations are supervised by the Spanish Agency for the Supervision of Artificial Intelligence. Breaching Article 50 can cost up to 15 million euros or 3 % of worldwide turnover.

- Article 50(1), disclose that a person is interacting with an AI, unless it is obvious, an exception read narrowly.
- Article 50(2), mark synthetic content in machine-readable form, text included, with a transitional period until 2 December 2026 for systems already on the market.
- Article 50(5), inform at the latest at the time of the first interaction, in a clear and distinguishable manner and meeting accessibility requirements.

The role this site plays

The assistant is offered under my own name even though the language model comes from a third party, so I take myself to be a provider under the regulation rather than merely a deployer. That is the reading with the heavier duties, and it is the one I have taken. Nor do I lean on the Article 2(10) carve-out for purely personal, non-professional use, because as I read it that covers natural persons deploying a system and not whoever puts it into service, and this site serves a professional purpose.

Measures in place

Each entry ties an obligation to what the site does and to the exact place in the code where it lives, which is what makes it auditable.

50(1) · Knowing you are talking to an AI

The chat panel is titled «José's AI assistant» and the greeting that opens the conversation says so again. Below the input there is a permanent notice reminding you that an AI is answering, that it can get things wrong and that the conversation is logged, with the privacy policy one click away.

```
src/components/chat/ChatPopover.tsx
```

50(1) · The assistant confirms it when asked

Its instructions require it to admit that it is an AI, in any phrasing and any language, above every other rule. What it still will not discuss is how it is built inside, which is a security matter rather than a transparency one.

```
src/lib/chat/prompts.ts
```

50(5) · In time, clear and accessible

The notice is on screen before the first message can be sent, and also in the quick search box, which sends the question before the panel opens. The text meets the WCAG 2.1 AA minimum contrast, which is what the law points to when it mentions accessibility.

```
src/components/chat/{ChatPopover,ChatSpotlight}.tsx
```

50(2) · Marking generated content

When a conversation or a fit analysis is downloaded or copied to be pasted elsewhere, the document carries a visible label and machine-readable metadata stating that an AI wrote it, when, and which system. It is the measure that strikes me as proportionate while no reliable watermark exists for free-form text.

```
src/lib/chat-transcript.ts
```

Guarding against drift

An automated test suite checks on every change that the notice is still there, that the assistant is still required to identify itself as an AI and that the exports are still marked. If a redesign takes any of them out, continuous integration fails before the deploy.

```
tests/unit/ai-transparency.test.ts
```

What does not apply, and why

Listing what falls outside matters as much as what falls inside, because it stops a future review from assuming something was simply overlooked.

- There are no deepfakes. When I reply to a voice or video message the recording is genuinely mine, so there is no synthetic content to label.
- There is no emotion recognition and no biometric categorisation, so as I read it Article 50(3) does not come into play.
- The assistant does not publish text intended to inform the public on matters of public interest, which is the case covered by the second subparagraph of Article 50(4).
- I do not consider it a high-risk system. The feature that analyses fit against a job posting brushes against Annex III, which mentions evaluating candidates, but the person using it is not the employer, the person evaluated is me, and the output neither filters nor ranks anybody's application. That regime is in any case deferred to December 2027.

What happens to the data

Article 50 transparency and data protection are different things and are documented separately. The detail of what is sent to the model, to which providers, on what legal basis and for how long it is kept lives in section 4 of this site's privacy policy, which opens from the chat itself. This report does not replace it.

When this gets reviewed

This document is revisited when the law or its interpretation changes, when the model or the gateway serving the assistant changes, and whenever a feature capable of generating new content is added. The nearest date still on the calendar is 2 December 2026, the end of the transitional period for machine-readable marking.

Sources

Regulation (EU) 2024/1689, Article 50

<https://artificialintelligenceact.eu/article/50/>

European Commission guidelines on transparency (20 July 2026)

<https://digital-strategy.ec.europa.eu/en/policies/guidelines-transparency-ai-generated-content>

Official FAQ on Article 50

<https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>

AESIA, code of practice on transparency

<https://aesia.digital.gob.es/es/actualidad/recursos/ria-articulo-50-faqs-codigo-de-buenas-practicas>

I am not a lawyer. This document describes technical measures and my own reading of the law, so it is neither legal advice nor a declaration of conformity. If anything it says does not match what you see on the site, write to me and I will correct it.

José Manuel Ortega de la Calle · Current version at www.josemanuelortega.dev/informe?lang=en