

EU AI Act Transparency Audit

What the providers of general-purpose AI models publish under the EU AI Act's transparency duties, read from their own public documents and checked against the date each duty falls due. Author: Mûcahit Kaya, AI Tools Police.

WHAT WE DID

We read the websites of the 26 providers on a list built in advance from two public platform rankings, between 7 October 2026 and 12 October 2026, and again on 14 October 2026. Each answer was coded from the provider's own documents under a protocol, sampling rules and codebook fixed and hash-recorded before collection and published with the dataset. For one model per provider we asked whether the summary of the content used to train it, which Article 53(1)(d) requires every provider of a general-purpose AI model to make publicly available, is published, whether anyone can open it without an account, and what it says, judged against its due date; and what providers publish about telling people they face an AI and about marking AI output. Every reading was AI-assisted, and 13 providers were read twice, blind. The products studied were not tried.

WHAT WE FOUND

11 of 21

EXPLORATORY: THE STUDY'S
FINAL CODING

Of the 21 providers whose summary deadline had passed for the model studied, **11 published a summary that presents itself as such and opens without an account (7 of 21 on the coding fixed before collection; Annex E).**

The final coding applies a broader identification rule than the one fixed before collection, and the summaries of Qwen and Tencent were reached only on 15 October 2026, after the collection window; Z.ai's summary names the GLM-5 family but not the model studied, and Qwen's describes the models behind its platform without naming one. Of the 2 providers whose deadline is 2 August 2027, 0 had published a summary.



- ☒ 11 summary that presents itself as such
- ☒ 4 related document only
- ☐ 6 nothing found in the places we looked

- **Related documents.** 4 of the same 21 published a related document describing training content without presenting itself as the summary: Google, inclusionAI, Sakana AI and Upstage. Related documents are counted on their own and never added to the summaries (Table 6).
- **Nothing found.** For 6 of 21 we found nothing in the places we looked: on the date of the read, in the kinds of page the codebook defines, we did not find the document. It does not mean the document does not exist (Table 5, Annex A).
- **The model matters.** These counts hold for the provider and model pairs studied: applied a second time, the rule that picks each provider's model reached the same model for 6 of 21 providers, and another model could change which document counts (Section 4.4).
- **The Commission's list.** In 9 places at a first read, and in 8 of them at 2 later rounds (11 September 2026, 29 September 2026 and 14 October 2026), we found no published list of general-purpose AI models with systemic risk, which the Act asks the Commission to publish. Article 52(6) sets no date by which the list must be published, and whether the time that has passed amounts to lateness is a question this study does not answer (Section 4.10).
- **Marking statements.** In the group whose marking deadline is 2 December 2026, 8 of 12 providers already state that outputs are marked; 9 of the 26 providers are removed from this question by a rule. Less certain: our two readings agreed below the thresholds set in advance on the 2 items this rests on (Section 4.11). We report what providers say about marking, not whether it works.

WHAT IT DOES NOT MEAN

- **No compliance or breach finding.** Nothing here is a finding that any provider, or the Commission, complied with or breached the law.
- **Documents only.** We read what providers publish on their own public pages, in English, following links. "Nothing found in the places we looked" is a result about those pages on that date.
- **The products were not tried.** The documents-only rule prohibited opening an account, filling in any form and using any model or product.
- **A constructed list, not a sample.** Nothing here supports a claim about providers outside the list of 26.

WHAT WE ASK

- **Methodological feedback** on the codebook, the identification rule and how the list was built.
- **Comparison with official guidance.** The Commission's questions and answers on the template ask that a summary covering several models or model versions specify which models and model versions it applies to. We would welcome that guidance applied to the summaries that do not name the version studied, Qwen's and Z.ai's, and a view on whether a related document can serve as the summary; this study does not decide either question (Sections 4.2 and 4.3).

WHERE THE DATA IS

The release (CC BY 4.0) holds every record: each answer's source address, reading date, quotation and the file name and SHA-256 of its saved copy; the copies themselves are not republished. **DOI:** [10.5281/zenodo.23415007](https://doi.org/10.5281/zenodo.23415007), <https://doi.org/10.5281/zenodo.23415007>. Contact: info@aitoolspolice.com