

EU AI Act Transparency Audit

A cross-sectional documentation audit of 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.

Wave 3. Collection window 7 October 2026 to 14 October 2026. Dataset frozen 14 October 2026. Final coding 16 October 2026.

Mücahit Kaya

AI Industry Researcher and Public Administration Specialist, AI Tools Police

ABOUT THIS REPORT

Author: Mücahit Kaya, AI Industry Researcher and Public Administration Specialist, AI Tools Police. **Study type:** Cross-sectional documentation audit of public provider and regulator materials. Observational. No intervention, no human subjects and no personal data; the products studied were not tried. **Protocol:** the research protocol, sampling rules and codebook, fixed and hash-recorded before collection and published with the dataset (AI Tools Police, 2026a; 2026b; 2026c). There was no public deposit before collection. **Status:** Preprint, 16 October 2026. Not peer reviewed.

CONTACT

Mücahit Kaya
info@aitoolspolice.com

HOW TO CITE THIS REPORT

Kaya, M. (2026) *EU AI Act Transparency Audit*. AI Tools Police. DOI: 10.5281/zenodo.23415007 (<https://doi.org/10.5281/zenodo.23415007>).

LICENSE

© AI Tools Police, 2026. The report and its dataset are published under the Creative Commons Attribution 4.0 International license (CC BY 4.0).

Independent research by AI Tools Police. This report is not a publication of the European Union or of any of its institutions, and it does not state their position.

Contents

1	Introduction	5
2	What the Act asks providers to publish	5
3	Method	6
4	Results	10
5	Limitations	20
6	Discussion	21
7	Recommendations: what would make the public record checkable (exploratory)	21
	Data availability	22
	Conflicts of interest, funding and relationships	24
	References	26
	Annex A. The training-content summary, provider by provider	29
	Annex B. Departures from the plan, in plain words	30
	Annex C. Glossary: plain labels and the codebook's codes	32
	Annex D. Consistency between the independent readings, item by item	38
	Annex E. The analysis on the coding fixed before collection, and the correction record	40
	AI-assistance disclosure	44

Abstract

Since 2 August 2025 the EU AI Act has required providers of general-purpose AI models to publish a summary of the content used to train them; models placed on the market before that date have until 2 August 2027. For one model per provider, we ask whether that summary 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.

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, coding each answer from the provider's own documents under a protocol fixed and hash-recorded before collection and published with the dataset. The products studied were not tried. Every reading was AI-assisted, and 13 providers were read twice, blind.

The central result, on the study's final, exploratory coding, is this. 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, so for neither can it be established from the document whether it applies to the version studied. Separately, 4 of the same 21 published a related document describing training content without presenting itself as the summary. Of the 2 providers whose deadline is 2 August 2027, 0 had published a summary. 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.

In the group whose marking deadline is 2 December 2026, 8 of 12 providers already state that outputs are marked (less certain: our two readings agreed below the thresholds set in advance on the 2 items this rests on; Section 4.11); 9 of the 26 providers are removed from this question by a rule. 26 of 26 publish a model card or other per-model documentation and 0 of 26 a copyright policy, neither of which the Act requires them to publish (Table 13). In the places we looked, on 3 dates, we found no published list of general-purpose AI models with systemic risk, which the Act asks the Commission to publish.

The two readings' consistency was 0.893 (Krippendorff's alpha) over every answer and 0.894 over the answers both readings coded directly; it is not a measure of accuracy. Under a name-matching rule fixed in advance, 21 of the 26 providers were matched to organizations in Epoch AI's database of AI models (20 once 3 name fragments are set aside); this compares two lists and validates neither. The list is a constructed frame, not a sample; nothing here finds that any company, or the Commission, complied with or breached the law.

KEY POINTS

- On the study's final, exploratory coding, of the 21 providers whose summary deadline had passed for the model studied, 11 published a training-content summary that presents itself as such and opens without an account (7 of 21 on the coding fixed before collection; Annex E). The summaries of Qwen and Tencent were reached on 15 October 2026, after the 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, so it cannot be established from either document whether it applies to the version studied.
- 4 more published only a related document, and for 6 we found nothing in the places we looked.
- The count belongs to the model studied for each provider: applied again, the rule that picks the model reached the same model for 6 of 21 providers.
- In the places we looked, we found no published list of general-purpose AI models with systemic risk.
- Nothing here is a finding that any provider, or the Commission, complied with or breached the law.

1 Introduction

The question is narrow, and it is the whole study: **for each provider studied, is the document the EU AI Act requires to be public actually published, can a person outside the company open it, and what does it contain, judged against the date from which it is due?**

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) set out its rules for general-purpose AI models, the large models many products are built on, in a chapter that has applied since 2 August 2025. Most duties in that chapter concern documents a provider gives to the authorities, or to other companies, on request. One is different. Article 53(1)(d) requires every provider of a general-purpose AI model to make a summary of the content used to train it **publicly available**. That makes the summary something any reader can look for, and so something a study can look for.

The question matters now for a plain reason. The Act's general date of application is 2 August 2026, under Article 113: "It shall apply from 2 August 2026." The earlier date for the chapter on general-purpose AI models, 2 August 2025, was set with one exception, written into the same article, Article 113, point (b): "Chapter III Section 4, Chapter V, Chapter VII and Chapter XII and Article 78 shall apply from 2 August 2025, with the exception of Article 101;" Article 101, which allows the Commission to fine providers of general-purpose AI models, was therefore held back to 2 August 2026, the same date from which the duties to tell people they are interacting with an AI system and to mark AI-generated content apply. This paper reads the public record at the moment those powers and duties have arrived; it makes no finding that bears on enforcement.

It is a question about documents, which is why documents are the data. A study that tried the products would answer a different question, and the products studied were not tried here. It reads what a provider publishes, records where it found each document and what the document says, and dates each finding. It does not judge whether any document is sufficient, and it does not say whether any company complies with the law.

Section 2 sets out what the Act asks, in plain terms. Section 3 describes the method in full. Section 4 reports the results, each under a heading that states the claim the section supports. Sections 5 and 6 give the limitations and what the results mean, and Section 7 the recommendations that follow from them. Every figure comes from the study's final analysis; Annex E gives the analysis of the coding as frozen and every value the final analysis corrects. The dataset release holds every record, every source and every table the protocol fixed; the Data availability section names each file.

2 What the Act asks providers to publish

Table 1 lists the provisions the study reads, whom each binds, whether the document must be public, and from when.

Table 1. The provisions this study reads.

Provision	What it asks (our summary)	Whom it binds	Must the document be published?	Applies from	Later deadline for earlier models or systems
Article 53(1)(d)	Draw up and publish a summary of the content used to train the model, using the AI Office template	Providers of general-purpose AI models	Yes	2 August 2025	Article 111(3): models on the market before 2 August 2025 have until 2 August 2027
Article 53(1)(a)	Keep technical documentation of the model, for the AI Office and national authorities on request	Providers of general-purpose AI models	No	2 August 2025	Article 111(3): models on the market before 2 August 2025 have until 2 August 2027
Article 53(1)(b)	Give information and documentation to developers who build the model into their systems	Providers of general-purpose AI models	No	2 August 2025	Article 111(3): models on the market before 2 August 2025 have until 2 August 2027
Article 53(1)(c)	Put in place a policy to comply with EU copyright law	Providers of general-purpose AI models	No	2 August 2025	Article 111(3): models on the market before 2 August 2025 have until 2 August 2027
Article 53(2)	Exempts certain open-source models from points (a) and (b)	No one (an exemption)	No	2 August 2025	None
Article 50(1)	Inform people that they are interacting with an AI system	Providers of AI systems	No	2 August 2026	Open: whether Article 111(3) reaches it (two readings reported)
Article 50(2)	Mark AI-generated outputs in a machine-readable, detectable way	Providers of AI systems	No	2 August 2026	Article 111(4): systems on the market before 2 August 2026 have until 2 December 2026
Article 50(4)	Disclose deep fakes and AI-generated public-interest text	Deployers	No	2 August 2026	None
Article 52(6)	The Commission ensures a list of general-purpose AI models with systemic risk is published	The Commission	Yes	2 August 2025	None

Note: our summary of each provision; whom it binds, the publication duty and the dates come from the study's frozen register of the Act's text, checked character for character against the consolidated text. Source: the obligation register in the release.

The training-content summary. Article 53(1)(d) requires providers to "draw up and make publicly available a sufficiently detailed summary about the content used for training of the general-purpose AI model, according to a template provided by the AI Office." The AI Office has published that template (European Commission, 2025a). The Act does not define "sufficiently detailed", and this study does not judge any summary against that phrase. Article 111(3) gives models already placed on the market more time: "Providers of general-purpose AI models that have been placed on the market before 2 August 2025 shall take the necessary steps in order to comply with the obligations laid down in this Regulation by 2 August 2027." We therefore dated each model by the earliest date on which its provider's own documents state it was available to outside users (Section 3.3). A model released before 2 August 2025 owes its summary by 2 August 2027; a model released on or after that date owes it from its release; where the documents did not settle the date, the provider is reported in a group of its own.

Documents the Act does not require to be public. Article 53(1) also requires technical documentation kept for the AI Office and national authorities, information for developers who build the model into their own systems, and a policy to comply with EU copyright law, including the rights reservation of Directive (EU) 2019/790. None of these has to be published. When providers publish them anyway, we count it as voluntary publication, with no deadline, and an absence is not a missing disclosure.

Interaction disclosure and marking. Article 50(1) requires AI systems that interact with people to be designed so that people are told they are dealing with an AI system. Article 50(2) requires providers of systems that generate audio, images, video or text to mark the outputs in a machine-readable way. Both apply from 2 August 2026. A later amendment, Regulation (EU) 2026/1744, added Article 111(4), which gives systems already placed on the market until 2 December 2026 to mark their outputs. Whether Article 111(3) also delays Article 50(1) for older models is an open question, so we report Article 50(1) under both readings and decide neither. We read what providers publish about these duties; we do not test whether any marking works. Article 50(4) binds deployers, the organizations that use a system, not providers, so guidance a provider publishes for its deployers is counted as voluntary.

The Commission's list. Article 52(6) requires the Commission to ensure that a list of general-purpose AI models with systemic risk is published and kept up to date. We looked for that list as carefully as we looked for any provider's document.

3 Method

3.1 Design and window

A cross-sectional documentation audit: one collection window, every provider on the frozen list read inside it, one instrument applied to all of them. Provider pages change, so one dated reading per provider is the unit of observation, and every answer carries its own reading date. Records were read between 7 October 2026 and 12 October 2026. Every record's documents were then read again on one day, 14 October 2026, in what this paper calls the final-day re-check; the plan put it on 13 October 2026, and it ran one day late (deviation 23). After the window, on 15 October 2026, an AI-assisted re-check read again the sources of the providers whose summary deadline had passed (Section 3.7). Every statistic is descriptive. The design answers no causal question, and no figure is a trend: this is one reading.

3.2 The frame: a constructed list, not a census or a sample

The providers were chosen in three steps, under rules fixed before any was read (AI Tools Police, 2026b). Table 2 gives the numbers.

Table 2. How the list of providers was built.

Step	Count
Candidates in the pool	176
Rejected by the eligibility tests	131
Candidates that passed	45
Providers those candidates make up	31
Providers with a position on at least one ranking (the list studied)	26
Eligible providers with a position on neither ranking	5
Providers with a system studied	19
Providers for whom no system qualified	7

Note: counts from the freeze record and the frozen list. Source: the freeze record and the analysis program's frame output.

First, candidates came from 7 public source pages: the 2 ranking platforms described below (Hugging Face, n.d.a; OpenRouter, n.d.a), 1 Commission page on the general-purpose AI code of practice (European Commission, n.d.e), and 4 catalog pages of 3 cloud services (Amazon Web Services, n.d.; Google, n.d.b; Microsoft, n.d.a; Microsoft, n.d.b); the pool closed on 29 September 2026. Second, eligibility tests ran from 29 September 2026 to 1 October 2026: a candidate entered only if its own English-language pages showed that it provides a general-purpose AI model, with a few logged cases decided on less than the candidate's own pages, or on a judgment. For 88 of the rejected candidates, the reason was that their own pages did not settle whether their model is general-purpose. Several candidates can belong to one company, which is why the candidates that passed make up fewer providers.

Third, the eligible providers were ranked on two public measures of how their models are distributed. One is the sum of the monthly download figures shown on the providers' model pages on Hugging Face (Hugging Face, n.d.b), read on 1 October 2026. The other is the number of tokens, the units of text a model processes, that the OpenRouter marketplace reports for each provider's models over a 7-day window, read on 1 October 2026 (OpenRouter, n.d.b). Neither measure shows size, revenue or importance: a download is a copy of a model file, not a user. Each provider took its better position; the rules admitted up to 50 providers, and every eligible provider with a position entered. The result is the 26 providers in this audit's ranked frame, frozen on 7 October 2026 before any document was coded.

The list is not neutral and supports no claim beyond itself. It is built from two platforms' figures, so it favors providers whose models are distributed through them, and it says nothing about providers below the cut. As a rough check on what a ranked list leaves out, the plan compares it with Epoch AI's public database of AI models (Epoch AI, 2026). Under the matching rule fixed before the comparison, 36 of the 1031 organizations in that database were matched to 21 of the 26 providers. 3 of those matches were fragments of company names, such as "Inc"; with them set aside, in a figure computed after the first was known, 33 of 1028 organizations match 20 providers (deviation 29). The rule matches names letter by letter and so cannot see a parent or brand name, such as Meta's or Alibaba's (deviation 30). The database describes itself as "The list of models is non-exhaustive, but aims to cover most models that were state-of-the-art when released, have over 1000 citations, one million monthly active users, or an equivalent level of historical significance." A provider not matched under the rule is not matched under it, and nothing more.

3.3 The instrument

The codebook is the measurement instrument (AI Tools Police, 2026c). It defines 48 coded items in nine areas: the training-content summary, the copyright policy, technical documentation and the open-source exemption, interaction disclosure, marking, guidance for deployers, what providers say about their own role under the Act, pages in other languages, and providers' own readings of terms of the Act. Each item has a closed list of answers, a decision rule, the kinds of page to read in order, and worked boundary cases. No item was added, removed or redefined after collection began, except that the final coding broadened the decision rule of the summary item, as the next paragraphs and Section 3.7 set out.

For each provider the frozen list fixed one model to study, which we call **the model studied**, and, where one qualified, one AI system offered to the public, such as a chat application, which we call **the system studied**. These choices fix which release date sets the deadline and which system the Article 50 questions are asked about. They are not claims that one model or system stands for everything a provider offers.

A document counts as **identified as the summary** when it presents itself as the summary of the content used to train the model. The codebook, fixed before collection, accepted four routes: the document names Article 53(1)(d), it names the AI Office template, its own title presents it as such a summary of a named model, or its opening sentence says so. The final coding applies one broader rule to every provider whose deadline had passed: a document presents itself as the summary where its title or its body says so, where it names Article 53(1)(d) or the AI Office template, or where it follows the template's section structure (deviation 36). A document that describes the model's training content but presents itself in none of these ways is a **related document**. Whether a document **names the model studied** is recorded apart: by its exact name, or not by name, where the document names only the model family or describes the provider's models in general; then whether it applies to the version studied cannot be established from it. Whether a related document, or a summary of a family, can be the summary the Act requires for a specific model is an open legal question, so related documents are counted on their own and never added to the summaries, and summaries that do not name the model studied are reported beside the count. The related documents of the providers outside the passed-deadline group, 2 whose deadline is 2 August 2027 and 1 whose release date was not settled, were coded under the codebook's routes only.

Every count about a duty is split by **deadline group**, computed from the provider's documented release date and the date we read the record: the deadline had passed; the deadline had not yet arrived; or the documents did not settle the release date. A group is a computation, not a judgment.

The release date is a proxy. The Act's duties attach to placing a model on the market in the Union, which a reader outside the company cannot observe. The study dates each model by one rule. The model's date is the earliest date on which the provider's own documents state that the model was available to users outside the company. A date called a "release date" that states no outside access does not qualify. Where two documents of the same standing give different dates, the date is recorded as conflicting, and the provider is reported in the group whose release date is not settled. The date so found is read as if it were the date of placing on the Union market. Where a model reached the Union after it was released elsewhere, the proxy dates it too early. That can put its provider in the passed-deadline group before the Act would, or, where the release elsewhere came before 2 August 2025 and the Union placing after it, in the group given until 2 August 2027 although the duty applied from the Union placing. Annex A prints the date the proxy takes for every model, with the deadline it sets. The second reading of Article 50(1), under which Article 111(3) delays that duty for older models, adds an assumption of its own: that the release date of the model studied stands for the system built on it. The Act states no such link, and the groups under that reading rest on it.

3.4 What was read, and what was prohibited

We read the provider's own public pages of nine kinds: the main website; legal, policy and trust pages; any page about the EU AI Act; model cards and other documentation of the model; developer documentation; the help center; the blog or newsroom; research papers the provider links; and the product page of the system. Each page was read in a real browser, and a copy of every page read was saved with its SHA-256 fingerprint. Pages behind a login or another barrier are reported apart from absences.

The documents-only rule prohibited, at every step: opening an account of any kind; signing in; filling in or submitting any form; starting a chat, a call or a purchase; typing into any search box on a provider's site; using any model or product; and contacting any provider or the Commission during the window. One breach is recorded: during the final-day re-check, a press of several controls on a Microsoft support page included controls that start a chat or a call. Nothing opened, nothing was typed, and no answer rests on that page (deviation 26).

We opened no search box on any provider's site. We found documents by following links, and we recorded separately whether each could be reached within three clicks of the home page. Within the window the study's only search was for the Commission's list, in the EU's law database; a general web search for the list was run at the first read, on 11 September 2026, before the window (Table 14). The AI-assisted re-check after the window also ran limited searches of providers' own domains, as its notes record. We read English-language pages.

"Nothing found in the places we looked" is a recorded result, and it means exactly that: on the date in that row, in those nine kinds of page, we did not find the document. It does not mean the document does not exist.

3.5 Who coded, and what the consistency figure measures

This study is AI-assisted research under named editorial control. One AI service carried out every reading in the window, first, second and adjudicating, under the documents-only rule, and a separate AI-assisted methods review read the plan, the collection rules and the paper adversarially. The AI-assisted re-check after the window, the two readings of deviation 38, the two readings and adjudication of deviation 39 and the two readings of deviation 40 are described in Section 3.7 and Annex B. The author fixed the question and the design, approved the protocol documents, and signed off the frozen dataset before publication.

The reliability figure is therefore reported as consistency between two independent readings, never as agreement between human coders. Two AI-assisted readings of the same pages, made by one service, can make the same mistake; where they do, they agree, the statistic rises, and accuracy does not. The figure shows that the instrument was applied consistently, not that it was applied correctly. Some answers are set by a codebook rule rather than coded, so two readings that agree on what triggers the rule agree on those answers automatically; the paper therefore also reports consistency over the answers both readings coded directly (Section 4.11). The protocol fixed the thresholds in advance: an item is flagged where Krippendorff's alpha, a standard agreement statistic (Krippendorff, 2019), falls below 0.667, or where the readings agree for fewer than 80% of providers.

3.6 Second coding, adjudication and blindness

A rule fixed in the protocol before any document was read chose 14 providers for a second, independent reading. Tencent left that group when its first reading was withdrawn twice, because collapsed answers on a read page could bear on the summary and had not been opened (deviation 22), so 13 providers were read twice. The second reading saw no answer of the first, no record of another provider and none of our own published pages, as each reading attested; the blindness was attested, not enforced by the system. Where the two disagreed, an adjudicating reading, blind on the same terms, read both records and both sets of sources and applied the codebook's decision rule. The published records come from the adjudicating reading for 11 providers and from the first reading for 15; among the latter, Moonshot AI and NVIDIA were read twice with no disagreement.

The blindness record. In 14 cases a reading reported that it had seen something outside its own provider's files, such as a list of file names, a count of files or the reading program's own source. A disclosure naming an answer, note, copy or page of the other reading of the same provider would have voided the reading; none did, so each reading stands. The 2 readings set aside were both Tencent's, for the reason above.

The final-day re-check and the re-readings. In the final-day re-check, on 14 October 2026, every record's documents were read again. Where it found a change that could bear on an answer, the record was read again in full: this happened for 4 providers, Anthropic, Cohere, Google and OpenAI. The study set the rule for these re-readings after their answers were known: a later reading replaces the earlier record only where a change the final-day re-check found in a provider's pages alters the answer it bears on (deviation 32). No change it found altered the answer it bore on, so the earlier records are published and the re-readings are kept beside them. The re-readings themselves differ from the published records on 12 answers, all listed in the release. The changes the final-day re-check found are reported as findings: Anthropic had rewritten the disclosure sentence of its acceptable use policy, and OpenAI's help article on the EU AI Act carried a new update date. For Cohere and Google it showed no change by the provider: one copy held only the site's navigation, and one showed content behind controls the earlier readings had not opened. On its own web crawling for training and its reading of what "as far as technically feasible" means for marking, Google's original record reads "Nothing found in the places we looked" while its re-reading found text; these 2 answers reflect the earlier reading, not an absence on Google's pages, and the final analysis applies the re-reading's values (deviation 34).

3.7 The original record, the correction record and the final analysis

The study keeps three records apart. The **original record** is the frozen records, one per provider, each answer as the reading coded it, signed off by the author before publication; the release's coded-value files carry it unchanged. The **correction record** lists every value the final analysis applies in place of an original one, with both values and the deviation each rests on: 37 entries, 35 of which change a value, for 7 providers (Table E4). Each applies a ruling of the deviations log to a published record (Cohere, deviation 33), takes the value the re-reading of 14 October 2026 found (Google, deviation 34), applies the broader identification rule to every provider in the group alike (deviation 36), counts a document our reading reached only after the window (deviations 36 and 37), codes the content of such a document from the study's own copy (deviation 38), or takes a value coded again by two blind readings (deviation 39 for Sakana AI's release date, deviation 40

for which models Qwen's summary names). Apart from Qwen's and Tencent's, whose summaries were read only after the window, none depends on what a provider's page showed after the window. The **final analysis** is the analysis program run over the original records with the correction record applied by program, and every figure in this paper comes from it unless the text says otherwise.

The headline item was re-coded after an external review. After an external review of an earlier version of this paper, and after every value was known, the item on whether a summary is published was coded again under the one rule set out in Section 3.3, applied to every provider whose deadline had passed (deviations 34 to 37). It is reported as the study's final, exploratory coding (protocol section 10.6), and the count on the coding fixed before collection stands beside it wherever it is the headline, and in Annex E. The re-coding rests on an AI-assisted re-check of the sources of the 21 providers whose deadline had passed, read on 15 October 2026. The AI-assisted re-check was not blind to the study's codes, was not time-limited, read after the window closed, and was not matched to the frozen copies; it answered apart whether each document presents itself as the summary and whether it names the model studied (Section 4.3). The author reviewed and confirmed its 21 rows on 15 October 2026; the confirmation is recorded in the study's log, not in the re-check file, whose review fields were left unchanged. Its readings changed the records of Microsoft, Qwen, Tencent and Thinking Machines Lab: those of Microsoft, Qwen and Thinking Machines Lab because the final coding applies the broader identification rule, not because an earlier reading was in error, and those of Qwen and Tencent because the re-check reached documents our reading had not; Qwen's move counts under both (deviation 37). Qwen's values rest on the study's own copy of its summary page, saved by its reader on 15 October 2026, and Tencent's on the study's own copy of its summary file, saved on 16 October 2026, identical by SHA-256 to the file the AI-assisted re-check opened. The AI-assisted re-check's own saved copies stay with the study and are not republished.

The content of the two summaries reached after the window. The content items of Tencent's and Qwen's summaries were coded from those copies by two readings under the codebook, each blind to the other, which agreed on 10 of 10 values, so no adjudication was needed (deviation 38). Two points of that coding are reported here once: Tencent's kinds of content rest on the boxes its summary ticks, which mark text only, and Qwen's named sources are two of its own models, named as generators of synthetic data.

3.8 The index and its sensitivity analyses

The study scores one thing, with an index the codebook names the Public Determinability Index for the Article 53(1)(d) summary: what a reader outside the company can establish about a document the Act requires to be public. It scores the public record, not the quality or adequacy of the content, and it scores no compliance. Its items and points, in Table 3, were fixed before any document was read. A provider with no document scores at the bottom of the scale, and a provider whose deadline has not arrived owes no summary yet, so a low band there is not a shortfall. The scores fall into five bands whose names describe the public record: full, substantial, partial, thin and minimal.

Table 3. The index items and their points.

Item	Maximum points
Training-content summary	25
The document names the model it covers	15
The document states the kinds of training content	10
The document states the categories of sources	15
The document names particular sources	10
The document refers to the AI Office template	10
The document carries a date or a version	15
Total	100

Note: each answer's points are fixed in the codebook; an answer the study could not determine scores zero and stays in the denominator, and the count of such answers is printed beside the index wherever it bears on it. Source: the analysis program's index output.

The weights are a fixed choice. The codebook set them before any document was read. Within an item it fixes the order of points: a published answer scores highest, a statement that the provider does not publish about a third of the item, documents that disagree about a fifth, and an answer that is absent or could not be settled scores nothing. Whether a summary is published and open carries the most points, and what it says shares the rest. The sensitivity analysis with equal weights shows how far the choice moves the index (Section 4.12).

Six sensitivity analyses were fixed in the protocol before collection. Two recompute the index: with equal weights for its items, and with every statement of non-publication scored zero. Four recompute the training-content figures for a subset or under another reading: signatories of the general-purpose AI code of practice, as the frozen list recorded them from Commission materials before any answer was coded; providers whose deadline had passed; the most favorable reading, under which an answer we could not open counts as found; and release dates assigned from a period, such as a season. None replaces a published figure or reorders a table (Section 4.12).

3.9 Departures from the plan, by class

The deviations log records 40 numbered departures and 4 notes, each dated, each stating what it changed, and the plan itself was never edited. Table 4 groups them; Annex B gives each in plain words.

Table 4. Departures from the plan, by class.

Class	Deviations	Count	Of which corrections of an earlier reading or rule
Building the list of providers	1 to 10	10	5
Rules for the readings, their blindness and what counts as a provider's document	11 to 22	12	2
The collection window, the final-day re-check and which record is published	23, 24, 25, 26, 27, 28, 32, 33	8	4
The outside comparison and the paper's own checks	29 to 31	3	0
The external review, the AI-assisted re-check and the final coding	34 to 40	7	2
All	1 to 40	40	13

Note: classes group the log's own entries by subject; the kind of each entry is the log's, except deviation 36, which deviation 37 describes as a re-coding under a broader identification rule rather than a correction. Source: the deviations log in the release.

Several departures were decided after the answers they bear on were known, and Annex B says so for each. These are the ones a reader may wish to open first. Deviations 1 to 4 corrected and widened the eligibility tests after the candidate pool and many decisions were known. Deviation 25 coded Z.ai's record again after its answers were known, and the summary still counts as identified (Section 4.3). Deviation 32 set the rule for the re-read records after their answers were known. Deviations 34 to 37 revised the analysis after reviews of an earlier version of this paper, deviation 36 re-coded the headline item, and deviation 38 coded the content of the two summaries reached after the window (Section 3.7 and Annex E).

Where the study holds two versions of a record, for the 4 providers read again and for Z.ai's record coded again, the analysis compares them answer by answer, and none of the differences changes which providers count in the headline figure. The re-coding of deviation 36 does change it, and Annex E sets the two analyses side by side. For the other departures the study holds no second version to compare, so their effect on the headline figure is not known.

3.10 Reporting rules

Fixed in the protocol before collection, these shape how this paper reads. Every share prints with its denominator, as "x of n"; percentages appear only for groups of 5 or more. Every count about a duty prints its deadline group, and no count pools two groups. Counts of documents the Act does not require to be public carry no deadline group. Providers a codebook rule removed from a question are printed beside the count as "removed by a rule" and leave its denominator. Every table about what providers publish carries the note "What is published and reachable, as read on the dates in each row. This study does not determine compliance." Every finding that rests on a flagged item carries a short caveat beside it. Per-provider tables are sorted by the study's provider identifier, never by score. No sentence says that any provider, or the Commission, complied with or breached the law.

4 Results

4.1 Where the deadline had passed, some providers publish a summary that presents itself as such, and for others we found nothing (exploratory)

This section reports the study's final, exploratory coding (protocol section 10.6), which applies the broader identification rule of Section 3.3 and counts documents reached after the window. The protocol fixes, word for word, a canonical sentence for the headline figure. Its words "at our read" would not hold for documents opened after the window, so it is printed in Annex E, filled on the coding fixed before collection, where every word of it holds.

On the final coding, of the 21 providers whose Article 53(1)(d) duty had already fallen due, 11 published a training-content summary that served live and that could be reached without an account (7 of 21 on the coding fixed before collection; Annex E). For 9 of them the document was read within the window, and for Microsoft and Thinking Machines Lab among them it counts as the summary only under the broader rule of the final coding; for Qwen and Tencent the document itself was reached on 15 October 2026, after the window. Of the 2 providers whose duty does not fall due until 2 August 2027, 0 published one, and for 3 the documents did not settle the release date. Table 5 breaks the same figures down by what we found, group by group.

Table 5. Training-content summaries, by deadline group.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

What we found	Deadline had passed (n = 21)	Not yet due, deadline 2 August 2027 (n = 2)	Release date not settled (n = 3)
Published, identified as the summary	11	0	0
Related document only	4	2	1
The provider states it does not publish this	0	0	1
Nothing found in the places we looked	6	0	1
Removed by a rule	0	0	0

Note: each cell counts providers; each column is one deadline group, with its own number of providers (n). Source: the analysis program's frequency output, final analysis.

The 11 providers with an identified summary, whose deadline had passed, are Anthropic, Cohere, DeepSeek, Microsoft, MiniMax, NVIDIA, OpenAI, Qwen, Tencent, Thinking Machines Lab and Z.ai. The summary names the model studied for 9 of them. For the other 2, 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, so for neither can it be established from the document whether it applies to the version studied (deviation 40). For 2 providers, Baidu and Meta, the deadline is 2 August 2027, and anything published now was published before it: 0 of the 2 published an identified summary, and 2 of the 2 a related document; whether such a document could serve as the summary the Act will require is the open question of Section 3.3. For 3 providers, Dolphin AI, Mistral AI and Salesforce, the documents did not settle the release date; 0 published an identified summary and 1 a related document.

Tencent's transparency center linked a training-content summary for its model studied, Hy4 preview, during the window (Tencent, n.d.b). The linked page was a PDF viewer that held no text, and the study's reading tool did not read the file (deviation 24 (c)). The AI-assisted re-check opened the file directly on 15 October 2026, and the study saved the same file on 16 October 2026 (Tencent, n.d.a). It is headed "Public Summary of Training Content for Hy4 Preview", and its file name dates it 28 August 2026. The final coding counts it as a summary that names the model studied, and its content was coded from the study's copy (deviation 38).

During the window, the footer of Qwen's home page showed the words "Training Data Summary" as plain text, with no link our reading recorded (Qwen, n.d.a), and the reading coded a summary mentioned with no location a reader can reach. After the window, on 15 October 2026, the study's own reader saved a training-content summary page on Qwen's site (Qwen, n.d.b); it describes the models of Qwen's platform without naming the model studied, Qwen3.8-Max. The Internet Archive's Wayback Machine index (Internet Archive, n.d.) lists the page's address with status 200 on 25 July 2026, 2 August 2026 and 26 August 2026, and on no date within the window: the address existed before the window, but no capture dates the page's content within it. The final coding counts it as a summary that describes the models behind Qwen's platform without naming one, by both the reading's later reach and the broader identification rule (deviations 36 and 37), and its content was coded from the study's copy (deviation 38).

Of the 15 providers in the passed-deadline group with a document address, 15 served live: Qwen's on 15 October 2026, and Tencent's on 15 October 2026 (opened by the AI-assisted re-check; the study saved the identical file on 16 October 2026); none was reached only through an archive copy. Of the 15 providers in that group whose summary or related document we located, the link crawl did not reach the document within three clicks of the home page for 2, and could not be completed, or was not run to a document reached after the window, for 9. Every provider's answer, with the document read and the release date that sets its deadline, is in Annex A.

4.2 Related documents describe training content without presenting themselves as the summary, and are counted apart

4 of the 21 providers whose deadline had passed published a related document, as did 2 of the 2 whose deadline is 2 August 2027 and 1 of the 3 whose release date was not settled. Table 6 lists them with what each document states.

Table 6. Related documents: what each states.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Company	Deadline group	Document read	Names the model it covers	Categories of sources	Named sources	Date or version
Baidu	Not yet due	(Baidu, n.d.)	Names specific models	Lists them	Gives examples of kinds of source only	Dated, no version
Google	Deadline had passed	(Google, n.d.a)	Names specific models	Not stated in the document	Names no source	Dated, no version
inclusionAI	Deadline had passed	(inclusionAI, n.d.)	Names specific models	Lists them	Names datasets	Dated and versioned
Meta	Not yet due	(Meta, n.d.)	Names specific models	Lists them	Names websites or licensors	Neither dated nor versioned
Sakana AI	Deadline had passed	(Sakana AI, n.d.)	Names specific models	Lists them	Gives examples of kinds of source only	Dated, no version
Salesforce	Release date not settled	(Salesforce, n.d.)	Names specific models	States them in general terms	Names datasets	Neither dated nor versioned
Upstage	Deadline had passed	(Upstage, n.d.)	Names specific models	States them in general terms	Names no source	Neither dated nor versioned

Note: sorted by the study's provider identifier. "Names the model it covers" records whether the document names specific models, not whether it names the model studied: inclusionAI's report names specific models but names the model studied only by its family (Table 7). Source: the final records.

Caveat: on the item "the document names particular sources", the two independent readings agreed for 10 of 13 providers (76.9%), below the 80% agreement threshold; Krippendorff's alpha was 0.694. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11).

Across summaries and related documents in the passed-deadline group, the document names the specific models it covers for 13 of 21 (61.9%) providers (for inclusionAI, not the model studied; Table 7) and lists the categories of sources for 13 of 21 (61.9%). It names particular datasets for 10 of 21 (47.6%) (less certain: our two readings agreed below a threshold set in advance on the item this rests on; Section 4.11). It names and cites the AI Office template for 1 of 21 (4.8%), and carries both a date and a version for 11 of 21 (52.4%). For 6 of 21 (28.6%) there was no document to read. These counts describe what the documents state, not whether it is sufficient.

4.3 Whether a document presents itself as the summary is kept apart from whether it names the model studied (exploratory)

The AI-assisted re-check of 15 October 2026 read the sources of every provider whose deadline had passed and answered two questions apart: whether the document presents itself as the summary, under the rule of Section 3.3, and whether it names the model studied by its exact name. Table 7 gives its reading for each provider, with the document it read and where. The final coding takes the first answer for every provider in the group; the second is reported beside it and changes no count.

Table 7. The AI-assisted re-check: does the document present itself as the summary, and does it name the model studied?

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Company	Document, as the AI-assisted re-check describes it	Where the AI-assisted re-check read it	Presents itself as the summary	Names the model studied	Final coding
Anthropic	Public Summary of Training Content: Claude Opus 5.5	The cover and section 1.2, PDF pages 1 and 2	Yes	Yes, by its exact name	Published, identified as the summary
ByteDance Seed	SeedRealtime model announcement	The announcement's heading and the model names in the Transparency list	No	Yes, by its exact name	Nothing found in the places we looked
Cohere	Training content summary for Command A+ and its family	PDF page 1, section 1.2	Yes	Yes, by its exact name	Published, identified as the summary
DeepSeek	DeepSeek-V4 training content summary	PDF page 1, section 1.2	Yes	Yes, by its exact name	Published, identified as the summary
Google	Gemini 3.8 Flash model card	The model card's description and training data sections	No	Yes, by its exact name	Related document only
inclusionAI	Ming-Flash-Omni technical report	PDF page 1 and the full text; also LEGAL.md	No	No; names its model family only	Related document only
Liquid AI	LFM2.5-230M product documentation	The model page opened from the model catalog	No	Yes, by its exact name	Nothing found in the places we looked

Company	Document, as the AI-assisted re-check describes it	Where the AI-assisted re-check read it	Presents itself as the summary	Names the model studied	Final coding
Meituan	LongCat changelog	The LongCat-2.5-Preview announcement; the terms of use	No	Yes, by its exact name	Nothing found in the places we looked
Microsoft	Data Summary for MAI-Thinking-1	PDF page 1, sections 1.1 to 1.3, and the structure of data sources that follows	Yes	Yes, by its exact name	Published, identified as the summary (originally: related document only)
MiniMax	MiniMax H3 training content summary	PDF page 1, the heading and section 1.2	Yes	Yes, by its exact name	Published, identified as the summary
Moonshot AI	Kimi K3 model card	The Hugging Face model card's heading and sections	No	Yes, by its exact name	Nothing found in the places we looked
NVIDIA	NVIDIA Nemotron training content summaries	PDF pages 18 and 19, the Nemotron 3.5 model list	Yes	Yes, by its exact name	Published, identified as the summary
OpenAI	GPT-6 Astra training content summary	PDF page 1, section 1.2	Yes	Yes, by its exact name	Published, identified as the summary
Perceptron	Perceptron Mk1.5 product documentation	The product document's heading and model identifier field	No	Yes, by its exact name	Nothing found in the places we looked
Qwen	Training Data Summary	The body of the summary, opened after the window from the link then at the foot of the home page	Yes	No; describes its platform's models without naming one	Published, identified as the summary (originally: summary mentioned, with no location a reader can reach)
Sakana AI	Sakana Fugu technical report	PDF page 1 and the full text; the Sakana Chat terms	No	Yes, by its exact name	Related document only
StepFun	Step 3.7 Flash model announcement	The model announcement's heading and the terms of use	No	Yes, by its exact name	Nothing found in the places we looked
Tencent	Hy4 Preview training content summary	PDF page 1, the heading and section 1.2	Yes	Yes, by its exact name	Published, identified as the summary (originally: could not be opened)
Thinking Machines Lab	Inkling-Small training content summary	PDF page 1, the heading and section 1.2	Yes	Yes, by its exact name	Published, identified as the summary (originally: related document only)
Upstage	Solar Pro 4 model announcement	The announcement's heading and the terms of use	No	Yes, by its exact name	Related document only
Z.ai	GLM-5 family training content summary	PDF page 1, section 1.2, and a search of the full text	Yes	No; names its model family only	Published, identified as the summary

Note: the descriptions and locators are the AI-assisted re-check's own, translated into English; a description is not a quotation of the document's title. The re-check did not code whether its documents describe training content, so the split between related documents and nothing found rests on the original reading. Source: the re-check output and the final records.

Under that rule, 9 of 21 documents present themselves as the summary and name the model studied, 2 of 21 present themselves as the summary but do not name the model studied, and 10 of 21 do not present themselves as the summary: 4 are the related documents of Table 5, and 6 are model cards, announcements or product pages; for these 6 the original reading found no document describing training content. The re-check finds a summary where the original coding counted none for 4 of 21 providers, Microsoft, Qwen, Tencent and Thinking Machines Lab, and finds none where the original coding counted one for 0 of 21. Microsoft's document (Microsoft, n.d.c) calls itself a data summary, names the model studied in its heading and follows the template's sections; Thinking Machines Lab's (Thinking Machines Lab, n.d.) names the model studied in its section on model identification. Z.ai's summary (Z.ai, n.d.) counted under both codings, as naming a model family: it names GLM-5, GLM-5.1, GLM-5.2 and GLM-5.3, not the model studied, GLM-5.3-Flash, and its date line reads "Last update: 21/08/2026", before that model's earliest documented date of outside access, 26 August 2026 (deviation 25). In its answer on modified or fine-tuned models, the Commission's questions and answers on the template say: "Different models or model versions may be covered by the same Summary, if the content of their Summaries is identical. In such cases, it should clearly specify which models and model versions the Summary applies to." (European Commission, 2026) Neither Qwen's nor Z.ai's summary specifies the model studied. The study does not judge whether either summary follows that guidance. The records of Microsoft, Tencent, Thinking Machines Lab and Z.ai, at the boundary the re-coding settled, were read once in the window; the AI-assisted re-check (Section 3.7) is the only other reading of them.

4.4 A different choice of model could change which documents count (exploratory)

The rule that picks each provider's model was applied a second time, independently, for 21 providers. It reached the same model for 6 of them. One of those, inclusionAI's, matches only once a hyphen and a space are read alike ("Ming-flash-omni-Preview" and "Ming-flash-omni Preview"), so letter for letter the same model was reached for 5 of 21; the two counts describe one comparison. Table 8 gives each pair.

Table 8. The model studied and the model the rule reached when applied again.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Company	Model studied	Model reached when the rule was applied again	Same model	Training-content summary, final coding	Model reached again named in that document
Anthropic	Claude Opus 5.5	Claude Sonnet 5.5	No	Published, identified as the summary	No
Baidu	ERNIE-4.5-0.3B	None reached	No	Related document only	No model reached
ByteDance Seed	SeedRealtime	SeedRealtime	Yes	Nothing found in the places we looked	Same model
Cohere	Command A+	Command A+	Yes	Published, identified as the summary	Same model
DeepSeek	DeepSeek-V4-Flash	DeepSeek-V4.1-Flash	No	Published, identified as the summary	No
Dolphin AI	Dolphin Mistral 24B Venice Edition	None reached	No	Nothing found in the places we looked	No model reached
Google	Gemini 3.8 Flash	Gemini 4 Argon	No	Related document only	No
inclusionAI	Ming-flash-omni-Preview	Ming-flash-omni Preview	Yes, once a hyphen and a space are read alike	Related document only	Same model
Liquid AI	LFM2.5-230M	LFM2.5-VL-3B	No	Nothing found in the places we looked	No document counted
Meta	Llama 4 Scout	Muse Spark 1.3	No	Related document only	No
MiniMax	MiniMax H3	MiniMax M3	No	Published, identified as the summary	No
Moonshot AI	Kimi K3	Kimi K2.5	No	Nothing found in the places we looked	No document counted
NVIDIA	NVIDIA-Nemotron-3.5-Lightning-30B-A3B-BF16	Nemotron 3 Nano 30B A3B	No	Published, identified as the summary	Yes
OpenAI	GPT-6 Astra	GPT-5.5	No	Published, identified as the summary	Yes
Perceptron	Perceptron Mk1.5	Isaac 0.5	No	Nothing found in the places we looked	No document counted
Qwen	Qwen3.8-Max	None reached	No	Published, identified as the summary	No model reached
Salesforce	Llama-xLAM-2-8b-fc-r	None reached	No	Related document only	No model reached
StepFun	Step 3.7 Flash	Step 3.7 Flash	Yes	Nothing found in the places we looked	Same model
Tencent	Hy4 preview	Hy4 preview	Yes	Published, identified as the summary	Same model
Thinking Machines Lab	Inklings-Small	Inklings-Small	Yes	Published, identified as the summary	Same model
Upstage	Solar Pro 4	None reached	No	Related document only	No model reached

Note: the models reached again were never coded, so no answer, deadline group or index exists for them; the last column tests only whether the document counted for the model studied also names the other model. Source: the analysis program's exploratory output on model selection, and the final records.

Among the 11 providers whose summary counts, the rule reached the same model again for Cohere, Tencent and Thinking Machines Lab. It reached a different model for Anthropic, DeepSeek and MiniMax, and their summaries do not name it, so on another selection those summaries might not have counted; for NVIDIA and OpenAI the other model is named in the same document. For Qwen the rule reached no model, and Microsoft and Z.ai were outside the group for which it was applied again. Because the models reached again were never coded, a different selection could also move a provider between deadline groups, by its release date, and this study cannot say in which direction. Every figure in this paper therefore describes the provider and model pairs on the frozen list, not each provider's models in general.

4.5 What a reader can establish about the summary varies widely between providers

Table 9 shows the index bands by deadline group. In the passed-deadline group the bands spread from full to minimal; a minimal band there records that we found no document in the places we looked.

Table 9. Index bands, by deadline group.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Band (index points)	Deadline had passed (n = 21)	Not yet due, deadline 2 August 2027 (n = 2)	Release date not settled (n = 3)
Full public record (85.0 to 100.0)	10 of 21 (47.6%)	0 of 2	0 of 3
Substantial public record (65.0 to 84.9)	1 of 21 (4.8%)	0 of 2	0 of 3
Partial public record (40.0 to 64.9)	2 of 21 (9.5%)	2 of 2	1 of 3
Thin public record (15.0 to 39.9)	2 of 21 (9.5%)	0 of 2	0 of 3
Minimal public record (0.0 to 14.9)	6 of 21 (28.6%)	0 of 2	2 of 3

Note: band names describe the public record, never a verdict; each column is one deadline group. Source: the analysis program's index output, final analysis.

Caveat: on the item "the document names particular sources", the two independent readings agreed for 10 of 13 providers (76.9%), below the 80% agreement threshold; Krippendorff's alpha was 0.694. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11). The index item on named sources rests on it.

Tencent's index is 90.0 (Full public record) and Qwen's 57.0 (Partial public record), computed on the content coded from the study's own copies of their summaries (deviation 38), and no index item of either is left not determinable. Anthropic's index is 90.0 (Full public record). Its model-report page was read with the tab for its model studied, Claude Opus 5.5, left unpressed, so the summary table under that tab was not read and the answers rest on the pages read; a sweep of every published record found this pattern at Anthropic alone (deviation 24 (b)).

4.6 Where a document sits, whether it opens, and what it says are reported apart (exploratory)

The index joins two things: whether the summary is published and open, which its first item scores, and what the summary says, which the other six score. Where the document sits on the site is recorded beside them and not scored. Table 10 reports the three apart for every provider, so that where a document sits, whether it opens and what it says can each be read on its own. The filling program computes it from the final analysis's item scores, and each row adds up to the index as the codebook computes it.

Table 10. The index in its parts: where the document sits, whether it is published and open, and what it says.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Company	Deadline group	Where the document sits (unscored)	Published and open (of 25)	Content (of 75)	Index (of 100), as the codebook computes it
Anthropic	Deadline had passed	One click from a page in the main menus	25	65	90.0
Baidu	Not yet due	Could not be determined (the link crawl could not be completed)	12	52	64.0
ByteDance Seed	Deadline had passed	No document to read	0	0	0.0
Cohere	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0
DeepSeek	Deadline had passed	Not reached within three clicks of the home page	25	65	90.0
Dolphin AI	Release date not settled	No document to read	0	0	0.0
Google	Deadline had passed	Three clicks from the home page, through standing pages	12	24	36.0
inclusionAI	Deadline had passed	Not reached within three clicks of the home page	12	65	77.0
Liquid AI	Deadline had passed	No document to read	0	0	0.0
Meituan	Deadline had passed	No document to read	0	0	0.0
Meta	Not yet due	Could not be determined (the link crawl could not be completed)	12	42	54.0
Microsoft	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0
MiniMax	Deadline had passed	Three clicks from the home page, through standing pages	25	68	93.0
Mistral AI	Release date not settled	No document to read	8	0	8.0

Company	Deadline group	Where the document sits (unscored)	Published and open (of 25)	Content (of 75)	Index (of 100), as the codebook computes it
Moonshot AI	Deadline had passed	No document to read	0	0	0.0
NVIDIA	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0
OpenAI	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0
Perceptron	Deadline had passed	No document to read	0	0	0.0
Qwen	Deadline had passed	Could not be determined (reached after the window; the crawl was not run to it)	25	32	57.0
Sakana AI	Deadline had passed	Could not be determined (the link crawl could not be completed)	12	42	54.0
Salesforce	Release date not settled	Could not be determined (the link crawl could not be completed)	12	32	44.0
StepFun	Deadline had passed	No document to read	0	0	0.0
Tencent	Deadline had passed	Could not be determined (reached after the window; the crawl was not run to it)	25	65	90.0
Thinking Machines Lab	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0
Upstage	Deadline had passed	Linked from the site's main menus	12	22	34.0
Z.ai	Deadline had passed	Could not be determined (the link crawl could not be completed)	25	65	90.0

Note: each row adds up to the index as the codebook computes it. Source: the analysis program's item scores and index, final analysis.

Caveat: on the item "the document names particular sources", the two independent readings agreed for 10 of 13 providers (76.9%), below the 80% agreement threshold; Krippendorff's alpha was 0.694. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11). The content points rest on it.

4.7 Statements that outputs are marked appear before the marking deadline

Table 11 reports what providers publish about marking their system's outputs as AI-generated, under Article 50(2). The question applies to providers whose system studied produces content; 9 providers are outside the table: 7 because no system qualified for study, and 2 because the documents read do not settle whether the system produces content (less certain: our two readings agreed below the thresholds set in advance on the 2 items this rests on; Section 4.11). The two cases are counted apart: of the 19 providers with a system studied, the provider's own documents state that it does not produce content for 0, and the documents read do not settle it for 2, 1 of them with a page our reading could not read.

Table 11. Statements that outputs are marked, by marking deadline group.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

What we found	Deadline had passed (n = 1)	Not yet due, deadline 2 December 2026 (n = 12)	System release date not settled (n = 4)
Outputs marked, form not stated	0	4	1
Outputs marked, in a machine-readable form	0	4	3
Marking planned	0	1	0
Nothing found in the places we looked	1	2	0
Could not be determined	0	1	0

Note: each column is one deadline group; the later deadline comes from Article 111(4). Source: the analysis program's frequency output.

Caveat: on the item "whether the provider's documents say the system produces content", the two independent readings agreed for 10 of 13 providers (76.9%), below the 80% agreement threshold; Krippendorff's alpha was 0.641, below the 0.667 threshold. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11). Which providers enter Table 11 depends on it.

Caveat: on the item "earliest documented public release date of the system", the two independent readings agreed for 9 of 13 providers (69.2%), below the 80% agreement threshold; Krippendorff's alpha was 0.647, below the 0.667 threshold. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11). The marking deadline groups rest on it.

For the 12 providers whose system studied was placed on the market before 2 August 2026, marking is due from 2 December 2026, and 8 of them already publish a statement that outputs are marked. Within that group, 4 of 12 (33.3%) name a marking specification or standard and 2 of 12 (16.7%) publish a route by which anyone can check an output (less certain: our two readings agreed below the thresholds set in advance on the 2 items this rests on; Section 4.11). The group whose deadline had passed holds Sakana AI, whose system studied, Sakana Chat, was first documented as public on 13 August 2026, so marking was due from that date. At our read on 8 October 2026 we did not find a statement that its outputs are marked, in the places we looked; undisclosed does not mean nonexistent (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.

4.8 On interaction disclosure, the reading of the deferral changes which providers are due

Article 50(1) applies from 2 August 2026. On the plain reading of its application date (reading A), the duty had arrived for every provider with a system studied. On the reading that Article 111(3) also delays it for models released before 2 August 2025 (reading B), the groups follow the release date of the model studied. Table 12 reports both; this study does not decide between them.

Table 12. Statements that users are told they are interacting with an AI system, under both readings.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Reading and group	Providers in group	Statement that users are told (x of n)	of which with the mechanism described	Nothing found in the places we looked	Could not be determined
Reading A (application date): already due	19	6 of 19 (31.6%)	4	12	1
Reading B (deferral): model released before 2 August 2025, due 2 August 2027	1	1 of 1	0	0	0
Reading B (deferral): model released on or after 2 August 2025, already due	15	5 of 15 (33.3%)	4	9	1
Reading B (deferral): model release date not settled	3	0 of 3	0	3	0
Reading B (deferral): in no group	0	No provider in this group			

Note: 7 providers have no system studied and are outside the table; reading B groups are dated by the model's documented release date. Source: the analysis program's frequency output.

On reading A, 6 of the 19 providers with a system studied publish a statement that users are told they are interacting with an AI system, and 4 of those statements describe the notice or the mechanism. On reading B, the duty of the 1 provider whose model was released before 2 August 2025, Meta, falls due on 2 August 2027; 1 of 1 already publishes such a statement. We report what providers say, not what a person using the system meets.

4.9 Model cards and developer documentation are widely published, and a published copyright policy was not found

Table 13 counts documents providers publish although the Act does not require them to. These counts carry no deadline group.

Table 13. Voluntary publication: documents the Act does not require providers to publish.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Document	Provision	To whom the Act directs it	Providers that published it	Removed by a rule
A copyright policy, presented as such	Article 53(1)(c)	Article 53(1)(c) names no recipient	0 of 26 (0.0%)	0
A statement that the technical documentation is kept (less certain, see caveat)	Article 53(1)(a)	The AI Office and the national competent authorities, on request	1 of 26 (3.8%)	0
A model card or other per-model documentation	Article 53(1)(a)	The AI Office and the national competent authorities, on request	26 of 26 (100.0%)	0
Documentation for developers who build on the model (less certain, see caveat)	Article 53(1)(b)	Providers of AI systems who intend to integrate the general-purpose AI model into their AI systems	24 of 26 (92.3%)	0
A license and the three public elements of the open-source exemption	Article 53(2)	Article 53(2) names no recipient	15 of 26 (57.7%)	0
Guidance for deployers on disclosing AI-generated content	Article 50(4)	Article 50(4) names no recipient for any document	10 of 19 (52.6%)	7

Note: the denominator is the providers to whom the question applies; for deployer guidance it is the providers with a system studied. A count of model cards is not a count of the technical documentation the Act directs to the authorities. Source: the analysis program's statements and frequency outputs.

Caveat: on the item "statement that technical documentation is kept for the authorities", the two independent readings agreed for 11 of 13 providers (84.6%); Krippendorff's alpha was 0.451, below the 0.667 threshold. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11).

Caveat: on the item "documentation for developers who build on the model", the two independent readings agreed for 12 of 13 providers (92.3%); Krippendorff's alpha was 0.648, below the 0.667 threshold. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11).

Article 53(1)(c) requires a copyright policy to be put in place and does not require it to be published, so the copyright count says nothing about whether any provider has one: 2 of 26 (7.7%) providers state that such a policy exists without publishing it. 10 of 26 (38.5%) publish a specific route by which a rightsholder can reserve their works against use in training, such as the name of the crawler to block. 6 of 26 (23.1%) name their own web crawler and say how to block it, and 9 of 26 (34.6%) state that they collected web content without naming a crawler. Google is among them in the final analysis, which applies the value its re-reading found (Section 3.6; deviation 34). Article 53(1)(c) names no recipient for any document, and no deadline attaches to these counts.

On what providers say about the Act, 5 of 26 (19.2%) state on their own pages that they provide a general-purpose AI model under it, and 12 of 26 (46.2%) do not mention it in the pages we read (less certain: our two readings agreed below a threshold set in advance on the item this rests on; Section 4.11). 10 of 26 (38.5%) publish a page about the EU AI Act. Where a provider states its own reading of a term of the Act, such as "sufficiently detailed", the record quotes it and grades nothing; every quotation is in the release.

4.10 We found no published list of models with systemic risk in the places we looked

We looked for the Commission's list in 9 places, in a real browser, at a first read and at 2 later rounds; the first read covered every place, and each later round every place except the general web search. Table 14 gives each.

Table 14. Where we looked for the Commission's list of models with systemic risk.

No.	Where we looked	First read, 11 September 2026	First round, 29 September 2026	Second round, 14 October 2026	Source
1	AI Act Service Desk page for Article 52	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.a)
2	Commission AI policy page (regulatory framework for AI)	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.g)
3	Commission questions and answers on general-purpose AI models	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.c)
4	AI Office page	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.d)
5	Direct address for a list of models with systemic risk	Address led to the site's page-not-found page; no list	Address led to the site's page-not-found page; no list	Address led to the site's page-not-found page; no list	(European Commission, n.d.f)
6	Commission guidelines on the scope of the general-purpose AI obligations	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, 2025b)
7	EUR-Lex search for the phrase "models with systemic risk"	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.h)
8	General web search for the list	Read; no published list found	Not read in this round	Not read in this round	Search; no single page
9	AI Act Service Desk page for Article 71 (EU database)	Read; no published list found	Read; no published list found	Read; no published list found	(European Commission, n.d.b)

Note: the later rounds do not repeat the general web search, because search engines met the study's automated browser with a human-verification challenge. Source: the study's record of the Commission checks and of the Act's primary texts.

We found no published list. 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. We also found no public register of the providers the Act's general-purpose AI duties bind: the EU database under Article 71 is limited, by its own text, to high-risk AI systems (European Commission, n.d.b), and the Commission's questions and answers on general-purpose AI models name no provider (European Commission, n.d.c).

4.11 The independent readings were consistent overall, and less so on the items that place systems in marking groups

Over every item together, consistency between the two independent readings was 0.893 (Krippendorff's alpha). Some answers are set by a codebook rule rather than coded: where no document exists, for instance, the questions about its content take the same answer. Leaving those out, over the 495 pairs of answers both readings coded directly, 443 of 495 agreed and alpha was 0.894, against 0.893 over all 621 pairs counted that way (exploratory; deviation 34). Neither figure measures accuracy (Section 3.5).

11 of the 48 items were flagged under the thresholds fixed in advance; Table 15 lists them with the rule that flagged each. No item was dropped for a low figure. Annex D gives all 48 items, with agreement, both alphas and the flag.

Table 15. The items flagged below the thresholds fixed in advance.

Item	Agreement before adjudication	Krippendorff's alpha	Rule that flagged it
What the provider says about its own role under the Act	10 of 13	0.677	agreement under 80%
Legal entity named by the provider	10 of 13	0.761	agreement under 80%
Earliest documented public release date of the system	9 of 13	0.647	agreement under 80% and alpha under 0.667
Whether the provider's documents say the system produces content	10 of 13	0.641	agreement under 80% and alpha under 0.667
The document names particular sources	10 of 13	0.694	agreement under 80%
Statement that technical documentation is kept for the authorities	11 of 13	0.451	alpha under 0.667
Documentation for developers who build on the model	12 of 13	0.648	alpha under 0.667
List of the kinds of output stated as marked	10 of 13	0.739	agreement under 80%
Own reading: a sufficiently detailed summary (Article 53(1)(d))	12 of 13	0.638	alpha under 0.667
Own reading: the open-source exemption (Article 53(2))	11 of 13	0.451	alpha under 0.667
Own reading: systemic risk (Article 51(2))	11 of 13	0.451	alpha under 0.667

Note: agreement is counted over the providers read twice, before adjudication. Source: the consistency figures file.

Of the two release dates that decide the deadline groups, the model's date was not flagged (alpha 0.92) and the system's date was (alpha 0.647): the training-content groups rest on the more consistent date and the marking groups on the less consistent one. The rule that picks each provider's system reached the same kind of system for 18 of 21 providers when applied again (Section 4.4 gives the same check for the model).

4.12 The sensitivity analyses fixed in advance replace no figure

Table 16 gives the result of each sensitivity analysis fixed in the protocol before collection.

Table 16. The sensitivity analyses fixed in the protocol.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Analysis fixed in advance	Figure it recomputes	Result
Index items weighted equally	The index	Differs by more than 5 points for 4 of 26 providers
Every statement of non-publication scored zero	The index	Differs by more than 5 points for 1 of 26 providers
Signatories of the code of practice only	Identified summaries, deadline passed	Signatories: 4 of 5 (80.0%); non-signatories: 7 of 16 (43.8%); signatory status not known: 0 providers; removed by a rule: 0 and 0
Providers whose deadline had passed only	Index bands	Full public record 10 of 21; Substantial public record 1 of 21; Partial public record 2 of 21; Thin public record 2 of 21; Minimal public record 6 of 21
Most favorable reading: answers we could not open count as found	Identified summaries, deadline passed	11 of 21 as published and 11 of 21 on this reading: no answer in the group is behind a barrier or could not be opened, so the figure does not move
Release dates assigned from a period, such as a season	Deadline groups	0 records carry such a date, so the check is untested

Note: none replaces a published figure or reorders a table. Source: the analysis program's sensitivity outputs, final analysis.

Caveat: on the item "the document names particular sources", the two independent readings agreed for 10 of 13 providers (76.9%), below the 80% agreement threshold; Krippendorff's alpha was 0.694. The thresholds were fixed in advance, so figures resting on this item are less certain (Section 4.11).

4.13 Answers the study could not settle, and to whom each kind belongs (exploratory)

Every coded answer that could not be settled is reported as itself and never counted as an absence. Table 17 groups them by item and says, for each group, to whom it belongs: our reading, the provider's documents, or neither. This grouping was not fixed in advance.

Table 17. Answers the study could not settle, by item and by whom they belong to.

Answer	Item	To whom it belongs	Answers	Providers
Could not be determined	A page in another language met during the reading	Our reading: a language switcher whose pages were not opened	13	13
Could not be determined	Where the document sits on the provider's site	Our reading: the link crawl could not be completed	10	10
Could not be determined	Other items	Our reading: pages referenced by a control with no address, or kinds of page not located or not read (deviations 21 and 33)	30	2
Could not be determined	Release date of the model or of the system studied	The provider's documents, which do not give the date	7	6
Could not be determined	Release date of a system, where no system was studied	No one: there is no system to date, and the item is coded although not asked in substance	7	7
Could not be determined	Where the summaries reached after the window sit on the provider's site	Our reading: the link crawl was not run to documents reached after the window (deviation 36)	2	2
The provider's documents disagree	Release date of the model studied	The provider's own documents, which give different dates	1	1

All answers asked: 1114 (1113 in the original record; the final coding also asks Qwen's question "summary or related document served live when read"); answers not settled: 70.

Note: answers a rule removed are not asked and not counted; values are those of the final analysis. Source: the final records.

Of the 70 answers that could not be settled, 55 belong to our reading. On pages in other languages, the codebook records "could not be determined" where a language switcher's pages were not opened, which was a choice of our reading. On where a document sits, it records the same where the link crawl could not be completed. The release dates are different: where a provider's documents give no date, the answer belongs to the documents, and where no system was studied there is no system to date at all. Where a kind of page was located and every read of it failed, an absence became "could not be determined" (deviation 21), which raised this count rather than the count of absences.

5 Limitations

This section lists what a reader should know before using any number here. The protocol's full register of limitations, item by item, is published with the dataset.

Not a compliance finding. "Nothing found in the places we looked" means one thing: on the date in that row, in the kinds of page the codebook defines, we did not find the document. No sentence here says that any provider or the Commission complied with or breached the law, and the study makes no legal assessment.

A constructed list, not a sample. The providers are those that two platforms' figures rank, after an eligibility test. Nothing here supports a claim about providers outside the list, and the list is small; the deadline groups make each figure smaller still.

One model and one system per provider, dated by the provider. A provider may publish many models, and every figure belongs to the model studied, not to the provider's models in general. The rule that picks the model reached the same model for 6 of 21 providers when applied again, and another model could change which document counts and which deadline group a provider sits in (Section 4.4). The deadline groups rest on release dates providers document about themselves, unevenly, and on a proxy: the earliest date on which the provider's own documents state the model was available to outside users, not placing on the Union market (Section 3.3).

AI-assisted readings, and a narrow consistency figure. One AI service made every reading in the window, and agreement between two of its readings is not accuracy. Blindness was attested by each reading, not enforced by the system. The AI-assisted re-check on which the final coding of the headline item rests was not blind and was made after the window (Section 3.7). The most useful check this design lacks is a person coding a subset against the same codebook, blind.

Where we looked. We read in English, from one location outside the European Union, following links rather than searching within the window; the AI-assisted re-check after the window also ran limited searches of providers' own domains. For 8 of 26 (30.8%) providers the reading met and recorded a page in another language, and for 13 of 26 (50.0%) whether it met one could not be determined. Some answers reflect our reading rather than a provider's silence; Section 4.13 counts them.

Decisions made after values were known. Several departures were decided after the answers they bear on were known (Annex B). The re-coding of the headline item after an external review is one of them, and it moves the headline count; Annex E gives both counts. Where the study holds two versions of a record from the re-readings, the headline count does not move; for the rest, the effect is not known. The re-coding applied its broader rule to the providers whose deadline had passed; the related documents of the other providers were not read again under it.

No publication dates, and one reading. The study cannot establish when any document was first published, measures no change over time, and promises no further reading. Two summaries were reached only after the window, and no capture dates the content of Qwen's page within it.

Commercial relationships. The study's AI service is provided by a company in the study, whose record was read under the same rules as every other provider's (Conflicts of interest).

6 Discussion

6.1 What the result means

The Act made one document about general-purpose AI models public by law, and that document can be looked for. On the study's final, exploratory coding, 11 of the 21 providers whose deadline had passed published a summary that presents itself as such; 9 of them name the model studied and 2 do not (7 of 21 on the coding fixed before collection; Annex E), and 4 of the same 21 published a related document that describes training content without presenting itself as the summary. From a provider's pages, a reader cannot always tell which document, if any, the provider means as the summary, or whether a summary covers the model the reader cares about.

On the regulator's side, in the places we looked, we found no published list of models with systemic risk and no public register of the providers the duties bind. A reader who wants to know which documents the Act requires of which company has to work that out from each company's own pages, as this study did.

6.2 How far the headline count can move, and why

The headline count carries three kinds of uncertainty, and no single range covers them all. **Classification:** on the coding as frozen the count was 7, and the final coding counts 11 (Annex E). The records of Microsoft, Qwen and Thinking Machines Lab moved because the final coding applies the broader rule to every provider in the group; which rule fits the Act is the open question of Section 3.3, and the AI-assisted re-check that applied it is described in Section 3.7. **Model selection:** applied again, the rule that picks the model reached the same model for 6 of 21 providers, and for Anthropic, DeepSeek and MiniMax the summary that counts does not name the other model (Section 4.4); the models reached again were never coded, so the effect cannot be counted. **Reading reach:** the records of Qwen and Tencent moved because the AI-assisted re-check reached documents our reading had not; Qwen's page was read after the window, and no capture dates its content within it; the content of both summaries was coded after the window, from the study's own copies (deviation 38); and 55 of the 70 answers the study could not settle belong to our reading rather than to the providers (Section 4.13). Each source is reported apart, beside the figure it bears on.

6.3 What the index is good for, and what it is not

The index is a determinability index: it measures what a reader can establish from the public record, not whether the record is adequate and not whether a provider complies. It separates providers with no document from providers with a dated, versioned summary that names its sources, and it is useful for that. It is not a ranking: it carries no claim about which provider is more transparent, and this paper prints no league table. A low band for a provider whose deadline has not arrived records a document not yet required.

6.4 What a further reading would need

If a further reading is made, the evidence here supports, in this order: a person coding a subset of providers against the same codebook, blind; the broader identification rule written into the codebook before collection, with coverage of the model coded apart; more than one model per provider, coded in full; a direct read of a PDF wherever a viewer page holds no text; a reading after 2 December 2026 and after 2 August 2027, when the later deadlines have arrived; and a second vantage point inside the European Union.

7 Recommendations: what would make the public record checkable (exploratory)

These are the author's recommendations, drawn from the findings above. They are not findings, and they are not part of the analysis the protocol fixed: they are exploratory, written after the values were known. Each is tied to the finding it answers. None is a statement that any provider or the Commission has complied with or breached the law. The question each answers is narrower: what would let a reader outside the company find the document, open it, and check what it says.

7.1 For providers

A summary that says what it is, and which model it covers. Article 53(1)(d) asks providers to "draw up and make publicly available a sufficiently detailed summary about the content used for training of the general-purpose AI model, according to a template provided by the AI Office." On the study's final, exploratory coding, 11 of the 21 providers whose deadline had passed published a summary that presents itself as such (7 of 21 on the coding fixed before collection; Annex E), and 4 of the same 21 published a related document that describes training content without presenting itself as the summary (Table 5). Under the codebook's narrower rule the documents of Microsoft, Qwen

and Thinking Machines Lab did not count as summaries, and under the broader rule of the final coding they do; 2 of the summaries do not name the model studied (Section 4.3). A reader could tell the summary apart at once if it were a standalone document whose own title is the template's title and which names, in its visible heading, every model version it covers.

A summary a reader can reach. Of the 15 providers in the passed-deadline group whose summary or related document we located, the link crawl did not reach the document within three clicks of the home page for 2, and could not be completed, or was not run, for 9 (Section 4.1). A link from a page in the site's main menus, or from a page about the EU AI Act that is itself linked from the main menus, would put the document within two clicks of the home page.

A summary that carries its date. In the same group, the document carries both a date and a version for 11 of 21 (52.4%) providers (Section 4.2). A date and a version number on the document itself let a reader tell which summary was current when, which one dated reading cannot establish.

A file that opens, at an address that lasts. Tencent linked a training-content summary, but the linked page was a PDF viewer that held no text at our read, and the file was opened directly only after the window (Section 4.1). A plain file that opens in any browser, at an address that does not change, lets any reader, and any public archive, keep a copy.

One place for the EU AI Act material. 10 of 26 (38.5%) providers publish a page about the EU AI Act (Section 4.9), and for 8 of 26 (30.8%) our reading met pages in another language, while for 13 of 26 (50.0%) whether it met one could not be determined (Section 5). One page, in English and in the provider's other languages, linking the summary, the model card and any statement about the code of practice, would give a reader one place to start.

7.2 For the European Commission and the AI Office

The list the Act requires. Article 52(6) reads: "The Commission shall ensure that a list of general-purpose AI models with systemic risk is published and shall keep that list up to date, without prejudice to the need to observe and protect intellectual property rights and confidential business information or trade secrets in accordance with Union and national law." In 9 places at a first read, and in 8 of them at 2 later rounds, we found no published list (Table 14). Publishing it would tell a reader which models carry the further duties for models with systemic risk.

One public index of summaries. We found no public register of the providers the duties bind (Section 4.10). An index kept by the AI Office, linking each provider's published summary with the model it covers and its date, would let a reader find every summary from one page, as this study had to do provider by provider.

The existing guidance, applied to the cases found. In its answer on modified or fine-tuned models, 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 (Section 4.3). Applying that guidance to the summaries found here that do not name the version studied, Qwen and Z.ai's, and saying whether a related document such as a model card with a section on training data can serve as the summary, would settle the cases on which the headline count turns. This study does not decide either question (Sections 4.2 and 4.3).

A machine-readable template. In the passed-deadline group, the document names and cites the AI Office template for 1 of 21 (4.8%) providers (Section 4.2). A machine-readable version of the template, with fixed field names, would let a reader, and a study, compare summaries field by field instead of reading each layout anew.

7.3 For researchers

What a repeat measurement would need. This study is one dated reading (Section 5). A repeat would need what this release provides: the file name and fingerprint of the saved copy behind every answer (the copies themselves are not republished); the address of every page; a link to a public archive capture where one exists, which the build found for 563 of the 1113 answers here; the codebook, the protocol and the sampling rules; and the analysis program and the filling program, so that every figure can be computed again. It would also need stable addresses from providers, which the recommendations above would supply.

Data availability

The release is deposited at Zenodo and published under a CC BY 4.0 license, DOI: 10.5281/zenodo.23415007

(<https://doi.org/10.5281/zenodo.23415007>). The documents the study read remain their providers' material. It contains these files, named as they stand in the study directory:

- `session-a-work/collection/frozen/*.json`: the original record, the frozen records, one per provider, each answer with its source address, reading date, quotation and the saved copy it rests on.
- `paper/v4/release/coded-dataset.csv`: the coded values in wide form, 26 rows, one per provider, in the column order the codebook fixes. **It holds the frozen values of the original record**; the final analysis applies the correction record to them, so a reader who counts this file reproduces Annex E, not the figures in the body of the paper.
- `paper/v4/release/evidence-long.csv`: the coded values in long form, 1113 rows, one per answer with its evidence, also holding the frozen values, in the eleven stored columns the codebook fixes and the two computed columns it adds: a link to a public archive capture

of each page, and the date of the build that found it. The build, run on 15 October 2026 with 238 queries of the Internet Archive's Wayback Machine index (Internet Archive, n.d.), linked a capture for 563 rows and found no capture that qualifies under the protocol's rule for 550, which say so; 0 rows name no copy. No value rests on either column: each answer rests on the study's saved copy.

- `analysis/out/final/corrections.csv`: the correction record, 37 entries, 35 of which change a value, each with its original value, the value the final analysis applies and the deviation it rests on.
- `paper/v4/release/correction-evidence.csv`: for every entry of the correction record, the source address, read date and saved copy, by name and SHA-256, on which the applied value rests.
- `session-a-work/collection/qwen-read-2026-10-15/wayback-cdx-answer.json`: the saved answer of the Wayback Machine index query for the address of Qwen's summary page.
- `session-a-work/collection/content-coding-2026-10-16/`: the two blind readings that coded the content of Tencent's and Qwen's summaries from the study's own copies of them (deviation 38).
- `session-a-work/collection/date-recoding-2026-10-16/`: the two blind readings and the adjudicating reading that coded the release dates of deviation 39 again.
- The Commission's questions and answers on the summary template, quoted in Section 4.3, rest on the study's copy `session-a-work/collection/ec-faq-read-2026-10-16/ec-faq-summary-template.html`, SHA-256 `e6afd4ac04c29ba2a9cacc4c09d443e177c6e544a34a16786baea9f99693614d`, which is the Commission's material and is not republished.
- `analysis/analyze.py` and `analysis/out/final/`: the analysis program and the final analysis, every output it writes over the original record with the correction record applied, with `analysis/out/final/MANIFEST.md` naming the rule each output implements and its SHA-256.
- `analysis/out/`: the analysis on the coding as frozen (Annex E), and the exploratory outputs: consistency over directly coded answers, model selection, the index in its parts on the coding before deviation 36, and the AI-assisted re-check's counts.
- `ai-assisted-recheck-2026-10-15/audit-results.csv`: the AI-assisted re-check's record of each provider whose deadline had passed, with the address, description and locator of the document it read, in Turkish as written; its saved copies are not republished.
- `codebook-v1.md`, `research-protocol-v1.md` and `sampling-rules-v1.md`: the instrument, protocol and sampling rules, fixed and hash-recorded before collection, the protocol carrying the full register of limitations.
- `obligation-register-2026-09-26.csv` and `primary-text-record.md`: the frozen register of the Act's text and the record of the primary texts read.
- `frame-frozen-2026-10-01.csv`, `frame-ranked-2026-10-01.csv`, `frame-rejected-2026-10-01.csv`, `frame-sources-2026-10-01.csv`, `crosscheck-2026-10-01.csv`, `double-coded-selection.csv`, `freeze-record-2026-10-01.md` and `freeze-checks-2026-10-01.csv`: the frame, its sources, the outside comparison and the freeze; the freeze record also holds the decisions under the registered control test that deviation 2 widened (Annex B).
- `collection-events-2026-10-07.csv`, `recheck-results.csv` and `session-a-work/collection/agree-2026-10-14.txt`: the events of collection, the Commission checks and the consistency figures.
- `deviations-log.md`: the departures from the plan.
- `paper/v4/paper-template.md`, `paper/v4/fill_paper.py`, `paper/v4/paper-trace.csv`, `paper/v4/build_release_tables.py` and `paper/v4/build_public_capture.py`: the paper's template, the program that fills every number, the trace naming the source of each, the program that builds the two coded-value files, and the program that builds the two public-capture columns. Section numbers are those of the PDF, whose renderer numbers the headings.
- `paper/v4/release/index-answers.json`: the archive index's answers the public-capture build read, so a rebuild gives the same columns without a new query.
- `paper/v4/release/checksums-sha256.txt`: the SHA-256 of every file above at the run that filled this paper, so a reader can confirm the copy held is the copy frozen.

The saved copies of provider pages, including the study's own copies of Qwen's and Tencent's summaries saved after the window, are not republished; each answer names its copy by file name and fingerprint, and the correction-evidence file does the same for every corrected value. The tables in this paper come from these files:

Table	Title	Built by the fill program from
Table 1	The provisions this study reads	<code>obligation-register-2026-09-26.csv</code>
Table 2	How the list of providers was built	<code>freeze-record-2026-10-01.md</code> ; <code>analysis/out/final/d13_frame_composition.csv</code>
Table 3	The index items and their points	<code>analysis/out/final/cb15_items_weights_bands.csv</code>
Table 4	Departures from the plan, by class	<code>deviations-log.md</code>
Table 5	Training-content summaries, by deadline group	<code>analysis/out/final/d01_frequencies.csv</code>

Table	Title	Built by the fill program from
Table 6	Related documents: what each states	session-a-work/collection/frozen/*.json
Table 7	The AI-assisted re-check: does the document present itself as the summary, and does it name the model studied?	analysis/out/recheck_d35.csv; session-a-work/collection/frozen/*.json; analysis/out/final/corrections.csv
Table 8	The model studied and the model the rule reached when applied again	analysis/out/model_selection_exposure.csv; session-a-work/collection/frozen/*.json
Table 9	Index bands, by deadline group	analysis/out/final/cb15_band_distribution.csv; analysis/out/final/cb15_items_weights_bands.md
Table 10	The index in its parts: where the document sits, whether it is published and open, and what it says	analysis/out/final/cb15_item_scores.csv; analysis/out/final/cb15_items_weights_bands.csv; analysis/out/final/cb15_index_pdi53d.csv; session-a-work/collection/frozen/*.json
Table 11	Statements that outputs are marked, by marking deadline group	analysis/out/final/d01_frequencies.csv
Table 12	Statements that users are told they are interacting with an AI system, under both readings	analysis/out/final/d01_frequencies.csv
Table 13	Voluntary publication: documents the Act does not require providers to publish	analysis/out/final/statement_traces.csv; obligation-register-2026-09-26.csv
Table 14	Where we looked for the Commission's list of models with systemic risk	recheck-results.csv; primary-text-record.md
Table 15	The items flagged below the thresholds fixed in advance	analysis/out/final/rp105_reliability.csv
Table 16	The sensitivity analyses fixed in the protocol	analysis/out/final/s104_pdi53d_equal_and_declared_zeroed.md; analysis/out/final/s104_p1_signatory_subset.csv; analysis/out/final/s104_p2_already_due_subset.md; analysis/out/final/s104_p3_most_favorable.csv; analysis/out/final/s104_p4_qualified_period.csv
Table 17	Answers the study could not settle, by item and by whom they belong to	session-a-work/collection/frozen/*.json; analysis/out/final/corrections.csv
Table A1	Training-content summary, provider by provider	session-a-work/collection/frozen/*.json; analysis/out/final/d19_due_state_membership.csv; analysis/out/final/corrections.csv
Table D1	Consistency between the two independent readings, every item	analysis/out/final/rp105_reliability.csv
Table E1	Training-content summaries, by deadline group, on the coding fixed before collection	analysis/out/d01_frequencies.csv
Table E2	The sensitivity analyses fixed in the protocol, on the coding fixed before collection	analysis/out/s104_pdi53d_equal_and_declared_zeroed.md; analysis/out/s104_p1_signatory_subset.csv; analysis/out/s104_p2_already_due_subset.md; analysis/out/s104_p3_most_favorable.csv; analysis/out/s104_p4_qualified_period.csv
Table E3	Every figure that differs between the analysis on the coding fixed before collection and the final analysis	analysis/out/d01_frequencies.csv; analysis/out/final/d01_frequencies.csv; analysis/out/d23_archive_only.csv; analysis/out/final/d23_archive_only.csv; analysis/out/d22_findability.csv; analysis/out/final/d22_findability.csv; analysis/out/cb15_band_distribution.csv; analysis/out/final/cb15_band_distribution.csv; analysis/out/cb15_index_pdi53d.csv; analysis/out/final/cb15_index_pdi53d.csv
Table E4	The correction record: every value the final analysis applies in place of an original one	analysis/out/final/corrections.csv

Conflicts of interest, funding and relationships

No external funding. No grant. No provider, platform or institution paid for this study or contributed to its design, and no person at any provider saw any part of it before publication. The study's AI service processed every reading in the window, including the reading of the record of its own maker.

AI Tools Police is reader-supported. The frozen list records, for each provider, our commercial relationships: affiliate relationships with 0 providers in the study and paid submissions from 0. Operational dependence: this publication depends operationally on Anthropic, a provider in this study, which also provides the study's AI service. Mitigations: the coding rules are written, with boundary cases reported

(Section 4.3); the second reading was blind to the first and the adjudication blind to our own published pages, as each reading attested; the index measures what a reader can establish, so a provider can raise it only by publishing clearer documents; and every answer publishes its source, its date and the name and fingerprint of its saved copy, so a reader can code a provider again. The saved copies of provider documents are named by file name and SHA-256 and are not republished.

References

AI Tools Police (2026a) 'Research protocol: EU AI Act Transparency Audit'. Study document fixed and hash-recorded before collection, published with the dataset.

AI Tools Police (2026b) 'Sampling rules: EU AI Act Transparency Audit'. Study document fixed and hash-recorded before collection, published with the dataset.

AI Tools Police (2026c) 'Codebook: EU AI Act Transparency Audit'. Study document fixed and hash-recorded before collection, published with the dataset.

AI Tools Police (n.d.) 'How we use AI'. Available at: <https://aitoolspolice.com/ai-transparency> (accessed 16 October 2026).

Amazon Web Services (n.d.) 'Models at a glance - Amazon Bedrock'. Available at: <https://docs.aws.amazon.com/bedrock/latest/userguide/model-cards.html> (accessed 29 September 2026).

Anthropic (n.d.) 'Public Summary of Training Content Claude Opus 5.5' [PDF; file title 'Claude Opus 5.5 Training Data Summary']. Available at: <https://trust.anthropic.com/doc/trust?rid=6ab2b1091d2552f21d552a3d&r=iz673w96495gyjer8h78n&download=true> (accessed 8 October 2026).

Baidu (n.d.) 'ERNIE 4.5 Technical Report' [PDF]. Available at: https://ernie.baidu.com/blog/publication/ERNIE_Technical_Report.pdf (accessed 7 October 2026).

Cohere (n.d.) 'Public Summary of Training Content for Command A+ and family' [PDF]. Available at: https://docs.cohere.com/_fern-files/cohere.docs.buildwithfern.com/465ad10da2a73b55ca4291e34e665bf843f8f918fd8a0c4e8f0092a8d60f93ef/assets/documents/eu-ai-public-summary_command-a-plus-family_20260910.pdf (accessed 8 October 2026).

DeepSeek (n.d.) 'Public Summary of Training Content for DeepSeek-V4' [PDF]. Available at: https://cdn.deepseek.com/policies/DeepSeek_Template_for_the_Public_Summary_of_Training_Content_for_GeneralPurpose_AI_model.pdf (accessed 8 October 2026).

Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, OJ L 130, 17.5.2019, p. 92.

Epoch AI (2026) 'Data on AI models', version 'Updated Oct. 5, 2026'. Available at: <https://epoch.ai/data/ai-models?view=graph&tab=notable> (accessed 7 October 2026).

European Commission (2025a) 'Explanatory Notice and Template for the Public Summary of Training Content for general-purpose AI models'. Published 24 July 2025. Available at: <https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models> (accessed 16 October 2026).

European Commission (2025b) 'Commission Guidelines on the scope of the obligations for providers of general-purpose AI models established by Regulation (EU) 2024/1689 (AI Act), C(2025) 7719 final'. Brussels, 19.11.2025. Available at: <https://ec.europa.eu/newsroom/dae/redirection/document/118340> (accessed 14 October 2026).

European Commission (2026) 'Template for general-purpose AI model providers to summarise their training content' [FAQ, last update 26 March 2026]. Available at: <https://digital-strategy.ec.europa.eu/en/faqs/template-general-purpose-ai-model-providers-summarise-their-training-content> (accessed 16 October 2026).

European Commission (n.d.a) 'Article 52: Procedure | AI Act Service Desk'. Available at: <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-52> (accessed 14 October 2026).

European Commission (n.d.b) 'Article 71: EU database for high-risk AI systems listed in Annex III | AI Act Service Desk'. Available at: <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-71> (accessed 14 October 2026).

European Commission (n.d.c) 'General-Purpose AI Models in the AI Act – Questions & Answers | Shaping Europe's digital future'. Available at: <https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers> (accessed 14 October 2026).

European Commission (n.d.d) 'European AI Office | Shaping Europe's digital future'. Available at: <https://digital-strategy.ec.europa.eu/en/policies/ai-office> (accessed 14 October 2026).

European Commission (n.d.e) 'The General-Purpose AI Code of Practice | Shaping Europe's digital future'. Available at: <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai> (accessed 29 September 2026).

European Commission (n.d.f) 'Page not found | Shaping Europe's digital future'. Available at: <https://digital-strategy.ec.europa.eu/en/policies/list-general-purpose-ai-models-systemic-risk> (accessed 14 October 2026).

European Commission (n.d.g) 'AI Act | Shaping Europe's digital future'. Available at: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai> (accessed 14 October 2026).

European Commission (n.d.h) 'Search results - EUR-Lex'. Available at: <https://eur-lex.europa.eu/search.html?text=%22models+with+systemic+risk%22&scope=EURL&type=quick&lang=en> (accessed 14 October 2026).

Google (n.d.a) 'Gemini 3.8 Flash - Model Card — Google DeepMind'. Available at: <https://deepmind.google/models/model-cards/gemini-3-8-flash/> (accessed 8 October 2026).

Google (n.d.b) 'Gemini Enterprise Agent Platform partner models for MaaS | Google Cloud Documentation'. Available at: <https://docs.cloud.google.com/gemini-enterprise-agent-platform/models/partner-models/use-partner-models> (accessed 29 September 2026).

Hugging Face (n.d.a) 'Models', sorted by downloads. Available at: <https://huggingface.co/models?sort=downloads> (accessed 29 September 2026).

Hugging Face (n.d.b) Model pages of the providers' models, each read for its monthly download figure; one address per model, listed in the frozen ranking file. Available at: <https://huggingface.co> (accessed 1 October 2026).

inclusionAI (n.d.) 'Ming-Flash-Omni: A Sparse, Unified Architecture for Multimodal Perception and Generation' [PDF]. Available at: <https://arxiv.org/pdf/2510.24821> (accessed 8 October 2026).

Internet Archive (n.d.) Wayback Machine index, queried for the address of Qwen's training-content summary page. Available at: <https://web.archive.org/cdx/search/cdx?url=qwen.ai/training-data-summary&output=json&fl=timestamp,original,statuscode,mimetype> (accessed 15 October 2026).

Krippendorff, K. (2019) Content Analysis: An Introduction to Its Methodology. 4th edn. Thousand Oaks, CA: SAGE.

Meta (n.d.) 'meta-llama/Llama-4-Scout-17B-16E-Instruct · Hugging Face'. Available at: <https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E-Instruct> (accessed 8 October 2026).

Microsoft (n.d.a) 'Foundry Models from partners and community - Microsoft Foundry | Microsoft Learn'. Available at: <https://learn.microsoft.com/en-us/azure/foundry/foundry-models/concepts/models-from-partners> (accessed 29 September 2026).

Microsoft (n.d.b) 'Foundry Models sold by Azure - Microsoft Foundry | Microsoft Learn'. Available at: <https://learn.microsoft.com/en-us/azure/foundry/foundry-models/concepts/models-sold-directly-by-azure?pivot=azure-openai> (accessed 29 September 2026).

Microsoft (n.d.c) 'Data Summary for MAI-Thinking-1' [PDF]. Available at: <https://microsoft.ai/pdf/MAI-Thinking-1-Data-Summary.pdf> (accessed 7 October 2026).

MiniMax (n.d.) 'Public Summary of Training Content for MiniMax H3' [PDF]. Available at: <https://file.cdn.minimax.io/public/8b640d09-235b-4351-a096-2f7bea762e75.pdf> (accessed 8 October 2026).

Mistral AI (n.d.) 'Does Mistral disclose its training datasets? | Mistral Help Center'. Available at: <https://help.mistral.ai/en/articles/347390-does-mistral-disclose-its-training-datasets> (accessed 7 October 2026).

NVIDIA (n.d.) 'Public Summary of Training Content for NVIDIA Nemotron Nano v2 12B VL' [PDF; a later part of the same file, which lists the model studied, is headed 'Public Summary of Training Content for NVIDIA Nemotron 3 and 3.5 Family']. Available at: <https://developer.download.nvidia.com/assets/nemo/docs/public-summary-of-training-content-for-nvidia-nemotron.pdf> (accessed 8 October 2026).

OpenAI (n.d.) 'Public Summary of Training Content for GPT-6 Astra' [PDF]. Available at: <https://cdn.openai.com/pdf/gpt-6-astra-eu-ai-act-public-summary-of-training-content.pdf> (accessed 8 October 2026).

OpenRouter (n.d.a) 'Models', ordered by weekly usage, read for the candidate pool. Available at: <https://openrouter.ai/models?order=top-weekly> (accessed 29 September 2026).

OpenRouter (n.d.b) 'Models', ordered by weekly usage, read for the token ranking. Available at: <https://openrouter.ai/models?order=top-weekly> (accessed 1 October 2026).

Qwen (n.d.a) 'Qwen'. Available at: <https://qwen.ai/home> (accessed 8 October 2026).

Qwen (n.d.b) 'Training Data Summary' [the page's heading; read by the study's reader after the window]. Available at: <https://qwen.ai/training-data-summary> (accessed 15 October 2026).

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024. Consolidated text as at 27 July 2026, CELEX 02024R1689-20260727, available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX%3A02024R1689-20260727> (accessed 24 September 2026).

Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24.7.2026. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ%3AL_202601744 (accessed 24 September 2026).

Sakana AI (n.d.) 'Sakana Fugu Technical Report' [PDF]. Available at: <https://arxiv.org/pdf/2606.21228> (accessed 8 October 2026).

Salesforce (n.d.) 'Salesforce/Llama-xLAM-2-8b-fc-r · Hugging Face'. Available at: <https://huggingface.co/Salesforce/Llama-xLAM-2-8b-fc-r> (accessed 8 October 2026).

Tencent (n.d.a) 'Public Summary of Training Content for Hy4 Preview' [PDF; saved by the study after the window]. Available at: <https://cdn-front.hy.tencent.ai/common-files/Hy4-Preview-Training-Data-Summary-20260828.pdf> (accessed 16 October 2026).

Tencent (n.d.b) 'Hy4-Preview-Training-Data-Summary-20260828'. Available at: <https://hy.tencent.ai/legal/Hy4-Preview-Training-Data-Summary-20260828> (accessed 12 October 2026).

Thinking Machines Lab (n.d.) 'Summary of Training Content for General-Purpose AI Model' [PDF; file title 'Inkling-Small - Public Summary of Training Content']. Available at: <https://thinkingmachines.ai/documents/inkling-small-public-summary-training-content.pdf> (accessed 7 October 2026).

Upstage (n.d.) 'Solar Pro 4: The Agentic Model That Finishes the Job'. Available at: <https://www.upstage.ai/blog/en/solar-pro-4> (accessed 8 October 2026).

Z.ai (n.d.) 'The Public Summary of Training Content for the GLM-5 Family' [PDF]. Available at: <https://cdn.bigmodel.cn/static/glm5.3/The%20Public%20Summary%20of%20Training%20Content%20for%20the%20GLM5%20Family.pdf> (accessed 9 October 2026).

Company documents and Commission web pages are cited without a year of publication, because this study does not establish when any document was published, except where a Commission document states its own date; the access date is the date of our read.

Annex A. The training-content summary, provider by provider

Table A1 gives each provider's model studied, the release date the deadline proxy takes for it, the deadline for its summary, what the final coding found and the document we read, so a reader can check any provider directly. The full record of every answer is in the release.

Table A1. Training-content summary, provider by provider.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Company	Model studied	Earliest date of outside access the provider's documents state (the deadline proxy)	Deadline for the summary	Training-content summary, final coding	Read on
Anthropic	Claude Opus 5.5	22 September 2026	Passed (due 22 September 2026)	Published, identified as the summary (Anthropic, n.d.)	8 October 2026
Baidu	ERNIE-4.5-0.3B	30 June 2025	Not yet due (due 2 August 2027)	Related document only (Baidu, n.d.)	7 October 2026
ByteDance Seed	SeedRealtime	5 August 2026	Passed (due 5 August 2026)	Nothing found in the places we looked	8 October 2026
Cohere	Command A+	20 May 2026	Passed (due 20 May 2026)	Published, identified as the summary (Cohere, n.d.)	8 October 2026
DeepSeek	DeepSeek-V4-Flash	24 April 2026	Passed (due 24 April 2026)	Published, identified as the summary (DeepSeek, n.d.)	8 October 2026
Dolphin AI	Dolphin Mistral 24B Venice Edition	Not settled	Release date not settled	Nothing found in the places we looked	7 October 2026
Google	Gemini 3.8 Flash	2 September 2026	Passed (due 2 September 2026)	Related document only (Google, n.d.a)	8 October 2026
inclusionAI	Ming-flash-omni-Preview	27 October 2025	Passed (due 27 October 2025)	Related document only (inclusionAI, n.d.)	8 October 2026
Liquid AI	LFM2.5-230M	25 June 2026	Passed (due 25 June 2026)	Nothing found in the places we looked	8 October 2026
Meituan	LongCat-2.5-Preview	25 September 2026	Passed (due 25 September 2026)	Nothing found in the places we looked	7 October 2026
Meta	Llama 4 Scout	5 April 2025	Not yet due (due 2 August 2027)	Related document only (Meta, n.d.)	8 October 2026
Microsoft	MAI-Thinking-1	12 August 2026	Passed (due 12 August 2026)	Published, identified as the summary (Microsoft, n.d.c); re-coded under deviation 36	7 October 2026
MiniMax	MiniMax H3	31 July 2026	Passed (due 31 July 2026)	Published, identified as the summary (MiniMax, n.d.)	8 October 2026
Mistral AI	Mistral Medium 3.5	Documents disagree	Release date not settled	The provider states it does not publish this (Mistral AI, n.d.)	7 October 2026
Moonshot AI	Kimi K3	16 July 2026	Passed (due 16 July 2026)	Nothing found in the places we looked	8 October 2026
NVIDIA	NVIDIA-Nemotron-3.5-Lightning-30B-A3B-BF16	11 August 2026	Passed (due 11 August 2026)	Published, identified as the summary (NVIDIA, n.d.)	8 October 2026
OpenAI	GPT-6 Astra	3 September 2026	Passed (due 3 September 2026)	Published, identified as the summary (OpenAI, n.d.)	8 October 2026
Perceptron	Perceptron Mk1.5	25 September 2026	Passed (due 25 September 2026)	Nothing found in the places we looked	7 October 2026
Qwen	Qwen3.8-Max	3 August 2026	Passed (due 3 August 2026)	Published, identified as the summary (Qwen, n.d.b); re-coded under deviation 36	8 October 2026; summary saved by the study 15 October 2026
Sakana AI	Fugu Ultra	24 April 2026 (re-coded under deviation 39)	Passed (due 24 April 2026)	Related document only (Sakana AI, n.d.)	8 October 2026
Salesforce	Llama-xLAM-2-8b-fc-r	Not settled	Release date not settled	Related document only (Salesforce, n.d.)	8 October 2026
StepFun	Step 3.7 Flash	29 May 2026	Passed (due 29 May 2026)	Nothing found in the places we looked	7 October 2026
Tencent	Hy4 preview	28 August 2026	Passed (due 28 August 2026)	Published, identified as the summary (Tencent, n.d.a); re-coded under deviation 36	12 October 2026; summary saved by the study 16 October 2026

Company	Model studied	Earliest date of outside access the provider's documents state (the deadline proxy)	Deadline for the summary	Training-content summary, final coding	Read on
Thinking Machines Lab	Inkling-Small	30 July 2026	Passed (due 30 July 2026)	Published, identified as the summary (Thinking Machines Lab, n.d.); re-coded under deviation 36	7 October 2026
Upstage	Solar Pro 4	6 August 2026	Passed (due 6 August 2026)	Related document only (Upstage, n.d.)	8 October 2026
Z.ai	GLM-5.3-Flash	26 August 2026	Passed (due 26 August 2026)	Published, identified as the summary (Z.ai, n.d.)	9 October 2026

Note: sorted by the study's provider identifier; the citation gives the document read for that answer; the release date is the earliest date on which the provider's own documents state the model was available to outside users, read as the date of placing on the Union market (Section 3.3). Source: the final records and the analysis program's deadline output.

Microsoft's summary (Microsoft, n.d.c) also states "Model release date: 2 June 2026" and "Date of placement of the model on the Union market: 12 August 2026". The study's date, 12 August 2026, is the first outside access Microsoft's documents record, a public preview, and it is also the date of placing on the Union market that the summary states.

Annex B. Departures from the plan, in plain words

Every departure from the plan fixed in the protocol was recorded in a dated public log during the work, and the plan itself was never edited. Each is listed here with its date and kind, what it changed and, where the log says so, that it was decided after the answers it bears on were known. The log gives the full text of each entry; where an entry names candidates, members or hosts, the log names them, and this annex names them only where the paper relies on them.

- **Deviation 1** (29 September 2026; Correction). An earlier reading of the eligibility test, which treated a product a candidate offers as a model it offers, was withdrawn. The decisions resting on it were made again; it reached candidates c151, c154, c161, c164, and Nomic and Nous Research, candidates seen after the pool closed, left the count.
- **Deviation 2** (29 September 2026; Departure from the plan). The test of whether a website belongs to a candidate was widened to follow the site's legal-notice link. It was decided after the candidate pool and 130 eligibility decisions were known, and among those decisions it changed candidates c004 and c153 only. Under the registered test both had been rejected on that control test (freeze record, lines 998 and 1673). Both are Google's, and Google is the only provider they bring into the frame: under the registered test the frame would hold the other 25 providers. That frame is not published as a separate file; the freeze record holds both decisions of each candidate.
- **Deviation 3** (30 September 2026; Correction). One standard for deciding whether a model is general-purpose was fixed and applied to every eligibility decision, after 136 decisions were known. Candidates it reached: c006, c117, c129, c153, c167, c169, c040, c043, c109, c113, c126, c128, c136, c147, c157, c165, c174, c014, c039, c049, c072, c076, c089, c004.
- **Deviation 4** (30 September 2026; Correction). Four readings of that standard were fixed by the methods review and adopted, after 202 decisions were known. Candidates it reached: c117, c097, c136, c018, c043, c064, c066, c071, c026.
- **Deviation 5** (1 October 2026; Departure from the plan). The token ranking's read was moved earlier than the sampling rules fixed, so its two reads fell closer together.
- **Deviation 6** (1 October 2026; Departure from the plan). That read was moved a second time, to the day the eligibility tests closed. The two token reads then shared most of their seven-day windows, and the download figures were read that day.
- **Deviation 7** (1 October 2026; Departure from the plan). The reads that chose each provider's model and system were bounded: at most two links from a starting page and sixteen page reads. A newer model beyond that bound could not be chosen. Which pages inside the bound were read was each reading's choice, and each absence means none on the pages read.
- **Deviation 8** (2 October 2026; Departure from the plan). For the choice of models and systems, the test of whether a website belongs to a provider was widened to other sites on the same domain and to platform accounts with a matching name. Ownership was then inferred from a domain or an account name that no page states.
- **Deviation 9** (2 October 2026; Correction). The choices of model for OpenAI, Qwen and Meta were corrected because the general-purpose test had been applied unevenly: OpenAI's model is now GPT-6 Astra, Qwen's Qwen3.8-Max and Meta's Llama 4 Scout.
- **Deviation 10** (7 October 2026; Correction). Our own review pages that set each provider's prior-coverage tier, which records how far our own site had reviewed the provider before the study, were widened to include our reviews of frontier models.
- **Deviation 11** (7 October 2026; A gap in the rules filled). A rule was set for a breach of the second reading's blindness: a disclosure naming an answer, note, copy or page of the other reading voids the reading; any other disclosure is listed and the reading stands. The readings set aside, 2 in all, were Tencent's, under deviation 22; none was set aside for a disclosure.
- **Deviation 12** (7 October 2026; Departure from the plan). The independent re-derivation of each provider's model and system was bounded as the original choice had been, starting from the provider's home page.
- **Deviation 13** (7 October 2026; A gap in the rules filled). The same rule was set for the adjudicating reading. No adjudication was voided.
- **Deviation 14** (7 October 2026; Departure from the plan). One ownership test was fixed for every reading, including single documents a provider's page links in its own name. Records already closed were coded again on their own saved copies, after records had closed;

each is in the release's events file.

- **Deviation 15** (7 October 2026; A gap in the rules filled). Cohere's first reading was withdrawn because one kind of page had failed without the third attempt and archive read the rules require. A new reading ran.
- **Deviation 16** (7 October 2026; Departure from the plan). A reading's note disagreeing with the choice of Mistral's model, which did not fit the note field, was kept whole beside the record. No answer moved.
- **Deviation 17** (8 October 2026; Correction, and a gap in the rules filled). Papers on other websites that a provider's pages link were admitted as the provider's research papers where they name the provider. Records it moved were coded again; Z.ai's reading was withdrawn and run again because a linked technical report had never been opened.
- **Deviation 18** (8 October 2026; A gap in the rules filled). A quotation from a PDF may run over several lines of one page. The rule was written after refusals on the records of Anthropic, Cohere and NVIDIA were known; records holding a PDF were coded again on their own copies.
- **Deviation 19** (8 October 2026; Departure from the plan). Where a reading saved a page but never wrote its read line, the line was completed from the reading tool's log. This was decided after it was known that the page could bear on the summary answer.
- **Deviation 20** (8 October 2026; Departure from the plan). The note route of deviation 16 was extended to links a reading left unopened (NVIDIA). It was entered after the answers were known and moved none.
- **Deviation 21** (8 October 2026; Correction, and a gap in the rules filled). An absence is a claim about all nine kinds of page, so where a kind of page was located but every read of it failed, "nothing found" became "could not be determined". It was applied after the answers it moved were known. The study computed no total of the values it moved; they are named in the dataset's events file, in the close event of each record coded again, and in the collection log.
- **Deviation 22** (8 October 2026; Departure from the plan). Tencent's first reading was withdrawn, twice, because collapsed answers on a read page could bear on the summary and had not been opened. This was decided after the first answer was known, and Tencent left the group read twice.
- **Deviation 23** (14 October 2026; Departure from the plan). The planned final day of collection, 13 October 2026, was missed, because the day was set by hand rather than by the protocol's rule that fixes it. The final-day re-check and the second round of the Commission checks ran one day late, on 14 October 2026. No answer of the final-day re-check was known when this was decided.
- **Deviation 24** (14 October 2026; Departure from the plan). Decided after the answers of the 4 re-readings were known. Part (a), which let the re-readings replace the earlier records, was later superseded by deviation 32. Part (b): Anthropic's model-report page was read with the tab for its model studied unpressed; the answers rest on the pages read (Section 4.5). Part (c): Tencent's summary could not be opened, and no answer moved. Part (d) concerned Cohere's page about the Act; once deviation 32 published Cohere's earlier record, deviation 33 applied the same ruling to that record's value.
- **Deviation 25** (14 October 2026; Correction). Z.ai's summary does not name the model studied, so its record was coded again on its own copies. This was decided after its answers were known. The summary still counts as identified; 3 answers moved.
- **Deviation 26** (14 October 2026; Departure from the plan). During Microsoft's final-day re-check, a press of several controls included controls that start a chat or a call, which the rules forbid. Nothing opened, nothing was typed or sent, and no answer rests on that page.
- **Deviation 27** (14 October 2026; Correction). Some rows of the second round of Commission reads named the reading tool's log file instead of the saved copy; the references were corrected to the copies.
- **Deviation 28** (14 October 2026; Correction). It corrects deviation 24, which had said the re-readings changed no answer. They did differ on other answers, and the release lists every one.
- **Deviation 29** (14 October 2026; Departure from the plan). The comparison with Epoch AI's database matched 3 name fragments, such as "Inc", to providers. The paper prints the comparison both with and without them, and the second figure was computed after the first was known.
- **Deviation 30** (14 October 2026; Departure from the plan). The same comparison cannot see a parent or brand name, so its figures are worded as matched or not matched under its rule.
- **Deviation 31** (14 October 2026; Departure from the plan). The quotation check now also reads curly quotation marks as straight ones and applies the whitespace rule for PDF text, and accepts no other difference.
- **Deviation 32** (14 October 2026; Departure from the plan). Set after the re-readings' answers were known: a re-reading made at the final-day re-check replaces the earlier record only where a change that re-check found in a provider's pages alters the answer it bears on. For the 4 re-read providers no such change altered the answer it bore on, so the study published the earlier records: the adjudicated record for Anthropic, Google and OpenAI, and the first reading for Cohere. The re-readings themselves differ from the published records on 12 answers, 2 of them Google answers the final analysis takes from the re-reading (deviation 34). The consistency figures now describe the published records of every provider read twice.
- **Deviation 33** (14 October 2026; Correction). With Cohere's earlier record published, its answer on a page about the Act rested on a link label with no address, which the ruling of deviation 21 counts as "Could not be determined". The analysis applies that value by program; the frozen record keeps the coded value, and the correction record lists both.
- **Deviation 34** (15 October 2026; Revision). Decided after an external review of an earlier version of this paper and after every value was known. The paper limits its claims to the provider and model pairs studied, states that the method documents were fixed and hash-recorded before collection and published with the dataset rather than deposited publicly, and states the proxy for the date a duty falls

due. The analysis program writes a corrected analysis beside the original, in which Google's two answers take the re-reading's values, and reports consistency over directly coded answers. The human audit this entry also planned was replaced by deviation 35.

- **Deviation 35** (15 October 2026; Addition). Decided after every value was known. In place of that human audit, an AI-assisted re-check of the sources of the providers whose deadline had passed, made after the window under the broader identification rule and not blind, is reported as exploratory; the author reviewed and confirmed its rows, as the study's log records. Its saved copies stay with the study.
- **Deviation 36** (15 October 2026; Re-coding under a broader identification rule, with two documents reached after the window). Decided after every value was known, superseding the side-by-side reporting of deviation 35. The confirmed readings of the AI-assisted re-check become the final coding of the headline item for Microsoft, Qwen, Tencent and Thinking Machines Lab, applied by program and listed in the correction record; the frozen records are unchanged, and the analysis on the coding as frozen is kept in Annex E. For Microsoft, Qwen and Thinking Machines Lab it applies a broader identification rule; it does not correct an error of the earlier reading. For Qwen and Tencent it counts documents our reading reached only after the window. It left the content of Tencent's and Qwen's summaries not determinable; deviation 38 coded it.
- **Deviation 37** (15 October 2026; Revision). Decided after a further review of an earlier version of this paper. Qwen's values rest on the study's own copy of its summary page, saved after the window, and the paper says that during the window the home page showed the label as plain text with no recorded link. The final coding is reported as exploratory, with the count on the coding fixed before collection beside it wherever it is the headline, and the index is printed as the codebook computes it, with the count of items not determinable beside it.
- **Deviation 38** (16 October 2026; Addition). Decided after every value was known. The content items of Tencent's and Qwen's summaries, left not determinable by deviation 36, were coded from the study's own saved copies by two readings, each blind to the other, which agreed on 10 of 10 values; the values are applied by program and listed in the correction record. Tencent's kinds of content rest on the boxes its summary ticks, which mark text only, and Qwen's named sources are two of its own models, named as generators of synthetic data.
- **Deviation 39** (16 October 2026; Correction). Decided after every value was known. The release date of the model studied was coded again for Microsoft, Mistral AI, Sakana AI and Tencent by two readings, each blind to the other, with an adjudicating reading where they differed. The dates of Microsoft, Mistral AI and Tencent stay as frozen. The date of Sakana AI becomes 24 April 2026, the first outside access to the model studied, in place of 11 September 2026, the date of a later version; it is applied by program and listed in the correction record. No deadline group changes.
- **Deviation 40** (16 October 2026; Correction). Decided after every value was known. The item on which models Qwen's summary names, set by deviation 36 without independent readings, was coded from the study's own copy of the summary by two readings, each blind to the other. Both found that the summary describes the models behind Qwen's platform without naming one; the value replaces the one deviation 36 set, is applied by program and is listed in the correction record. The summary still counts as identified, and no deadline group or count of summaries changes; Qwen's index moves with the item's points.

The log also holds 4 notes on these entries. One records the exact time of the decision behind deviation 29. One adds that the quotation check also treats special space characters as spaces (deviation 31). One records that the note deviation 24 (b) requires could not be written beside Anthropic's final record and is printed in this paper instead (deviation 32). One completes deviation 36: Qwen's summary counts as served live, read after the window, and the answers on where Tencent's and Qwen's summaries sit become "Could not be determined", because the link crawl was not run to them.

Annex C. Glossary: plain labels and the codebook's codes

This annex is the only place where the codebook's variable and value codes appear; Annex B names candidates by their identifiers in the frame files. It maps every plain label this paper prints to the codebook's variable and value, so that a methods reader can check any figure against the dataset. The program that fills this paper prints Tables C1 to C5 from the same dictionary it uses everywhere else.

Table C1. Items and their values.

Plain label (item)	Codebook variable	Plain label of a value: codebook value
Training-content summary	tcs_published	
		Published, identified as the summary: identified_summary_public
		Related document only: adjacent_document_public
		Summary behind a login or another barrier: identified_summary_gated
		Related document behind a login or another barrier: adjacent_document_gated
		Summary mentioned, with no location a reader can reach: referenced_no_location

Plain label (item)	Codebook variable	Plain label of a value: codebook value
Summary or related document served live when read	tcs_live_serve_status	Served live when read: served_live
		Reached only in an archive copy: archive_only
		Not retrieved by our reading, live or from an archive copy: neither
Where the summary or related document sits on the provider's site	tcs_findability	Linked from the site's main menus: nav_named_document
		One click from a page in the main menus: one_click_from_nav_named
		Three clicks from the home page, through standing pages: deeper_in_site_hierarchy
		Reached only through a dated post: single_post_link
		Not reached within three clicks of the home page: not_found_in_bounded_crawl
		No document to read: none_stated
The document names the model it covers	tcs_identifies_model	Could not be determined (the link crawl could not be completed): unknown
		Names specific models: names_specific_models
		Names a model family: names_model_family
		Covers the provider in general: covers_provider_generally
The document states the kinds of training content	tcs_modalities_stated	No document to read: none_stated
		Lists them: stated_enumerated
		States them in general terms: stated_general
The document states the categories of sources	tcs_source_categories_stated	No document to read: none_stated
		Lists them: stated_enumerated
		States them in general terms: stated_general
The document names particular sources	tcs_named_sources_stated	No document to read: none_stated
		Names datasets: named_datasets
		Names websites or licensors: named_domains_or_licensors
		Gives examples of kinds of source only: class_examples_only
		Names no source: none_named
The document refers to the AI Office template	tcs_template_reference	No document to read: none_stated
		Names and cites the template: template_named_and_cited
		Mentions the template without citing it: template_mentioned_without_citation
		No reference to the template: no_template_reference
The document carries a date or a version	tcs_currency	No document to read: none_stated
		No document to read: none_stated
Copyright policy	copyright_policy_public	

Plain label (item)	Codebook variable	Plain label of a value: codebook value
		Policy published, open without an account: <code>policy_document_public</code>
		Policy said to exist, not published: <code>policy_described_not_published</code>
		Statements about copyright, no policy mentioned: <code>copyright_statement_only</code>
The copyright policy refers to EU copyright law	<code>copyright_policy_eu_law_referenced</code>	
		Cites the EU text and data mining reservation: <code>dsm_article_4_3_cited</code>
		Refers to EU copyright law: <code>union_copyright_law_referenced</code>
		Refers to copyright in general: <code>copyright_referenced_generally</code>
		No policy document to read, or no reference: <code>none_stated</code>
Route for rightsholders to reserve their rights against training	<code>rights_reservation_mechanism_documented</code>	
		A specific public route (crawler name, opt-out form or signal): <code>route_public_and_specific</code>
		A route described without specifics: <code>route_described_no_location</code>
Statement on the provider's own web crawling for training	<code>tcs_scraping_disclosed</code>	
		Crawler named, with how to block it: <code>crawler_named_and_exclusion_documented</code>
		Crawler named only: <code>crawler_named_only</code>
		Web collection stated, crawler not named: <code>scraping_stated_crawler_unnamed</code>
The copyright policy carries a date or a version	<code>copyright_policy_currency</code>	
		No policy document to read: <code>none_stated</code>
Statement that technical documentation is kept for the authorities	<code>techdoc_regulator_route_documented</code>	
		Statement, with a route for the authorities: <code>statement_and_route_public</code>
		Statement only: <code>statement_only</code>
Documentation for developers who build on the model	<code>downstream_integrator_documentation_public</code>	
		Behind an account: <code>account_required</code>
		Available on request: <code>on_request_route_documented</code>
Model card or other per-model documentation	<code>model_card_public</code>	
The model card carries a date or a version	<code>model_card_currency</code>	
Conditions of the open-source exemption stated	<code>oss_exemption_conditions_documented</code>	
		License and all three public elements stated: <code>all_elements_documented</code>
		Some of the elements stated: <code>some_elements_documented</code>
Statement on whether the model has systemic risk	<code>systemic_risk_designation_stated</code>	
		States the model has systemic risk: <code>states_systemic_risk_model</code>
		States the model does not have systemic risk: <code>states_not_systemic_risk</code>
		States a training-compute figure, no designation: <code>compute_figure_stated_no_designation</code>

Plain label (item)	Codebook variable	Plain label of a value: codebook value
Statement that users are told they are interacting with an AI system	interaction_disclosure_documented	
		Statement, with the notice or mechanism described: mechanism_documented
		Statement, no mechanism described: stated_without_mechanism
		States no notice is given because the AI nature is obvious: obviousness_relied_on
That statement carries a date or a version	interaction_disclosure_currency	
Whether the provider's documents say the system produces content	system_generative_status	
		Produces content, in explicit words: generation_stated_explicit
		Produces content, in ordinary product language: generation_stated_functional
		States it does not produce content: generation_stated_negative
		The documents read do not settle it: generation_unknown
The coded generative status agrees with the system's selection	designation_status_agreement	
		Agrees: agrees
Statement that the system's outputs are marked	marking_disclosure_published	
		Outputs marked, in a machine-readable form: marking_stated_machine_readable
		Outputs marked, form not stated: marking_stated_form_unspecified
		Marking planned: marking_stated_planned
The marking method is named	marking_method_named	
		Names a specification or standard: named_specification
		Describes a method, names no specification: method_described_unnamed
		No method stated: unstated
		No marking statement: none_stated
The kinds of output the marking covers	marking_modalities_covered	
		Lists the kinds of output marked: modalities_enumerated
		States that all outputs are marked: all_outputs_stated
		Not stated: unstated
		No marking statement: none_stated
List of the kinds of output stated as marked	marking_modalities_list	
Which models carry the marking, and from when	marking_scope_dated	
		Models and start date stated: scope_models_and_date_stated
		Models stated, no date: scope_models_stated_undated
		Date stated, models not: date_stated_models_unspecified
		Neither stated: unstated
		No marking statement: none_stated
Public route to check an output for the marking	marking_verification_route_documented	
		A public route to check outputs: public_route_documented
		A route behind an account: route_documented_account_required

Plain label (item)	Codebook variable	Plain label of a value: codebook value
		A route described, with no location: route_described_no_location
The marking statement carries a date or a version	marking_disclosure_currency	
		No marking statement: none_stated
Guidance for deployers on disclosing AI-generated content	deployer_guidance_published	
		Guidance citing Article 50: guidance_cites_article_50
		Guidance without a reference to Article 50: guidance_without_article_reference
What the provider says about its own role under the Act	provider_gpai_self_identification	
		States it provides a general-purpose AI model under the Act: states_gpai_provider
		Mentions the Act, states no role: references_act_without_role
		States it is not such a provider: declared_absent
		Does not mention the Act: no_reference_to_act
A page about the EU AI Act	eu_act_resource_page_public	
Statement on adhering to the code of practice	code_of_practice_signature_stated	
		States adherence, names chapters: signature_stated_chapters_named
		States adherence, names no chapter: signature_stated_chapters_unnamed
		States it does not adhere: declared_absent
Legal entity named by the provider	provider_legal_entity_stated	
How the model is released (open or closed weights)	release_form	
		Open weights: open_weight_stated
		Closed weights: closed_weight_stated
		Mixed: mixed_stated
		Not stated: unstated
Earliest documented date of outside access to the model	model_public_availability_first_documented_date	
Earliest documented public release date of the system	system_public_availability_first_documented_date	
A page in another language met during the reading	non_english_surface_encountered	
		Met and recorded: encountered_recorded
		None met: none_encountered
Own reading: when the AI nature of a system is obvious (Article 50(1))	interp_quote_50_1_obviousness	
Own reading: marking as far as technically feasible (Article 50(2))	interp_quote_50_2_feasibility	
Own reading: assistive editing functions (Article 50(2))	interp_quote_50_2_carveout	
Own reading: the exemptions and limits in Article 50(4)	interp_quote_50_4_scope	
Own reading: a sufficiently detailed summary (Article 53(1)(d))	interp_quote_53_1d_sufficiency	
Own reading: the AI Office template (Article 53(1)(d))	interp_quote_53_1d_template	
Own reading: the open-source exemption (Article 53(2))	interp_quote_53_2_exemption	

Plain label (item)	Codebook variable	Plain label of a value: codebook value
Own reading: systemic risk (Article 51(2))	interp_quote_51_2_systemic_risk	

Note: the plain labels the paper prints, mapped to the codebook's variable names and values.

Table C2. Values shared by many items.

Plain label	Codebook value (any item, unless the item maps it otherwise above)
Nothing found in the places we looked	none_stated
Could not be determined	unknown
Could not be opened	unretrievable
Behind a login or another barrier	gated
The provider's documents disagree	conflicting
The provider states it does not publish this	declared_absent
Not asked, because a rule removed the question	not_applicable
Not stated in the document	unstated
No provider in this group	(empty_group)
Dated and versioned	dated_and_versioned
Dated, no version	dated_only
Versioned, no date	versioned_only
Neither dated nor versioned	undated_unversioned
Published, open without an account	public_url
Published, open without an account	public_page
Mentioned, with no location a reader can reach	referenced_no_location

Table C3. Deadline groups, readings and index bands.

Plain label	Codebook token
Deadline had passed (training-content summary)	due_at_collection
Not yet due (training-content summary)	not_due_at_collection
Release date not settled (training-content summary)	not_determinable
Deadline had passed (marking)	due_at_collection
Not yet due (marking)	not_due_at_collection
System release date not settled (marking)	not_determinable
Reading A (application date): already due	due_at_collection (every record on which the items are coded)
Reading B (deferral): model released before 2 August 2025, due 2 August 2027	designated model available before 2 August 2025: not due until 2 August 2027
Reading B (deferral): model released on or after 2 August 2025, already due	designated model available on or after 2 August 2025: falls due as on the application-date reading
Reading B (deferral): model release date not settled	documents did not settle when the designated model became publicly available
Reading B (deferral): in no group	beside the three groups (D19 item 7): art53_stratum not_applicable, in no group
Full public record	full_public_record
Substantial public record	substantial_public_record
Partial public record	partial_public_record
Thin public record	thin_public_record
Minimal public record	minimal_public_record
Not computed	not_applicable

Table C4. The rules that remove or answer a question.

Codebook dependency rule	What it does, in plain words
D1a	no summary or related document, so the questions about its content and place are answered from that
D1b	the summary question could not be settled, so its dependent questions take the same answer
D1c	the summary could not be opened, so its dependent questions take the same answer
D2	no marking statement, so the questions about the marking are answered from that
D3	no system qualified for study, so the questions about a system are not asked
D4	the system is documented as not producing content, so the marking questions are not asked
D4b	the documents do not settle whether the system produces content, so the marking questions are not asked
D4c	the coded generative status contradicts the system's selection, so the marking questions are not asked
D5	no copyright policy document, so the questions about it are answered from that
D6	no model card, so its date question is answered from that
D7	the provider was not admitted as a provider of an AI system, so the system questions are not asked
D7b	the provider was not admitted as a provider of a general-purpose AI model
D8	no document address exists to test, so the live-serving question is not asked

Note: a provider "removed by a rule" is one for whom one of these rules took the question away; it leaves that item's denominator.

Table C5. Events of the collection and reasons for rejecting a candidate.

Plain label	Code
Reading withdrawn	attempt_withdrawn
Record coded again on its own copies	record_given_again
Reading set aside	pass_voided
Reader's note kept beside the record	note_line_returned
Missing read line completed	read_line_completed
Record read again in full on the day of the final-day re-check	recollection
Reading abandoned	pass_abandoned
Difference noted at the final-day re-check	recheck_difference
Line range of a quotation from a PDF recorded	quote_lines
All events	all_events
Rejected: not a provider of AI models on its own pages	g1_not_a_model_provider
Rejected: its own pages did not settle whether its model is general-purpose	g2_not_established
Rejected: its own pages describe its models as built for one task or a narrow set	g2_not_general_purpose
Rejected: not the party providing the model	g3_not_the_provider
Rejected: no English-language page of its own that a reader could reach	g6_not_first_party_reachable

Note: the event types of the events file and the reason codes of the eligibility tests.

Annex D. Consistency between the independent readings, item by item

The protocol requires the consistency statistics to be published as their own table, whatever the values are. Table D1 gives every coded item with its agreement before adjudication, its nominal alpha, its secondary alpha where the codebook names one, and whether it fell below a threshold fixed in advance.

Table D1. Consistency between the two independent readings, every item.

Item	Agreement before adjudication	Krippendorff's alpha (nominal)	Secondary alpha (n)	Below threshold
What the provider says about its own role under the Act	10 of 13	0.677	ordinal 0.664 (n 13)	Yes
A page about the EU AI Act	11 of 13	0.749	ordinal 0.588 (n 13)	No
Statement on adhering to the code of practice	11 of 13	0.797	ordinal 0.958 (n 13)	No

Item	Agreement before adjudication	Krippendorff's alpha (nominal)	Secondary alpha (n)	Below threshold
Legal entity named by the provider	10 of 13	0.761	None	Yes
How the model is released (open or closed weights)	12 of 13	0.862	ordinal 0.978 (n 13)	No
Earliest documented date of outside access to the model	12 of 13	0.92	ratio 1.0 (n 12; dates as day numbers)	No
Earliest documented public release date of the system	9 of 13	0.647	ratio 0.964 (n 7; dates as day numbers)	Yes
Whether the provider's documents say the system produces content	10 of 13	0.641	ordinal 0.714 (n 13)	Yes
The coded generative status agrees with the system's selection	13 of 13	1.0	None	No
Training-content summary	11 of 13	0.791	ordinal 0.883 (n 13)	No
Summary or related document served live when read	12 of 13	0.837	ordinal 0.837 (n 13)	No
Where the summary or related document sits on the provider's site	11 of 13	0.799	ordinal 0.764 (n 13)	No
The document names the model it covers	12 of 13	0.837	ordinal 0.837 (n 13)	No
The document states the kinds of training content	11 of 13	0.789	ordinal 0.894 (n 13)	No
The document states the categories of sources	12 of 13	0.885	ordinal 0.981 (n 13)	No
The document names particular sources	10 of 13	0.694	ordinal 0.923 (n 13)	Yes
The document refers to the AI Office template	12 of 13	0.863	ordinal 0.877 (n 13)	No
The document carries a date or a version	11 of 13	0.789	ordinal 0.882 (n 13)	No
Copyright policy	13 of 13	1.0	ordinal 1.0 (n 13)	No
The copyright policy refers to EU copyright law	13 of 13	1.0	ordinal 1.0 (n 13)	No
Route for rightsholders to reserve their rights against training	12 of 13	0.874	ordinal 0.949 (n 13)	No
Statement on the provider's own web crawling for training	12 of 13	0.897	ordinal 0.974 (n 13)	No
The copyright policy carries a date or a version	13 of 13	1.0	ordinal 1.0 (n 13)	No
Statement that technical documentation is kept for the authorities	11 of 13	0.451	ordinal 0.518 (n 13)	Yes
Documentation for developers who build on the model	12 of 13	0.648	ordinal 0.581 (n 13)	Yes
Model card or other per-model documentation	13 of 13	Not defined (no variation)	ordinal undefined (n 13)	No
The model card carries a date or a version	13 of 13	1.0	ordinal 1.0 (n 13)	No
Conditions of the open-source exemption stated	12 of 13	0.874	ordinal 0.949 (n 13)	No
Statement on whether the model has systemic risk	12 of 13	0.775	ordinal 0.799 (n 13)	No
Statement that users are told they are interacting with an AI system	12 of 13	0.89	ordinal 0.935 (n 12)	No
That statement carries a date or a version	12 of 13	0.89	ordinal 0.938 (n 13)	No
Statement that the system's outputs are marked	12 of 13	0.905	ordinal 0.989 (n 13)	No
The marking method is named	12 of 13	0.908	ordinal 0.989 (n 13)	No
The kinds of output the marking covers	11 of 13	0.8	ordinal 0.963 (n 13)	No
List of the kinds of output stated as marked	10 of 13	0.739	None	Yes
Which models carry the marking, and from when	11 of 13	0.812	ordinal 0.964 (n 13)	No
Public route to check an output for the marking	12 of 13	0.893	ordinal 0.964 (n 13)	No
The marking statement carries a date or a version	12 of 13	0.909	ordinal 0.975 (n 13)	No
Guidance for deployers on disclosing AI-generated content	13 of 13	1.0	ordinal 1.0 (n 13)	No
Own reading: when the AI nature of a system is obvious (Article 50(1))	12 of 13	0.775	None	No

Item	Agreement before adjudication	Krippendorff's alpha (nominal)	Secondary alpha (n)	Below threshold
Own reading: marking as far as technically feasible (Article 50(2))	11 of 13	0.701	None	No
Own reading: assistive editing functions (Article 50(2))	12 of 13	0.832	None	No
Own reading: the exemptions and limits in Article 50(4)	12 of 13	0.775	None	No
Own reading: a sufficiently detailed summary (Article 53(1)(d))	12 of 13	0.638	None	Yes
Own reading: the AI Office template (Article 53(1)(d))	12 of 13	0.775	None	No
Own reading: the open-source exemption (Article 53(2))	11 of 13	0.451	None	Yes
Own reading: systemic risk (Article 51(2))	11 of 13	0.451	None	Yes
A page in another language met during the reading	12 of 13	0.828	None	No

Note: agreement is counted over the providers read twice; the secondary alpha treats ordered answers in order, or dates as day numbers. Source: the consistency figures file.

Annex E. The analysis on the coding fixed before collection, and the correction record

The final analysis applies the correction record to the original record; the analysis on the coding as frozen, which the protocol fixed, is kept unchanged beside it (Section 3.7). This annex prints its headline results whole. The protocol's canonical sentence, filled by the analysis program on the coding fixed before collection, reads:

Of the 21 providers in this audit's ranked frame whose Article 53(1)(d) duty had already fallen due at our read, 7 published a training-content summary that served live and that we could reach without an account. Of the 2 providers whose duty does not fall due until 2 August 2027 under Article 111(3), 0 published one. For 3 providers the documents did not settle when the designated model became publicly available, and those are reported separately. A further 0 providers are not in any of those three groups, because the dependency rules removed the Article 53(1)(d) variables from their records.

In that sentence, "the designated model" is the model studied. Table E1 is Table 5 on that coding. Table E2 gives the six sensitivity analyses on it. Table E3 lists every figure that differs between the two analyses, found by program, and Table E4 the correction record whole.

Table E1. Training-content summaries, by deadline group, on the coding fixed before collection.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

What we found	Deadline had passed (n = 21)	Not yet due, deadline 2 August 2027 (n = 2)	Release date not settled (n = 3)
Published, identified as the summary	7	0	0
Related document only	6	2	1
Summary mentioned, with no location a reader can reach	1	0	0
The provider states it does not publish this	0	0	1
Nothing found in the places we looked	6	0	1
Could not be opened	1	0	0
Removed by a rule	0	0	0

Note: each cell counts providers; each column is one deadline group. Source: the analysis program's frequency output on the frozen coding.

Table E2. The sensitivity analyses fixed in the protocol, on the coding fixed before collection.

What is published and reachable, as read on the dates in each row. This study does not determine compliance.

Analysis fixed in advance	Figure it recomputes	Result
Index items weighted equally	The index	Differs by more than 5 points for 3 of 26 providers
Every statement of non-publication scored zero	The index	Differs by more than 5 points for 1 of 26 providers
Signatories of the code of practice only	Identified summaries, deadline passed	Signatories: 3 of 5 (60.0%); non-signatories: 4 of 16 (25.0%); signatory status not known: 0 providers; removed by a rule: 0 and 0
Providers whose deadline had passed only	Index bands	Full public record 7 of 21; Substantial public record 3 of 21; Partial public record 1 of 21; Thin public record 2 of 21; Minimal public record 8 of 21
Most favorable reading: answers we could not open count as found	Identified summaries, deadline passed	7 of 21 as published; 8 of 21 on this reading, counting 1 answer that could not be opened
Release dates assigned from a period, such as a season	Deadline groups	0 records carry such a date, so the check is untested

Note: none replaces a published figure or reorders a table. The index rows rest on a flagged item (less certain: our two readings agreed below a threshold set in advance on the item this rests on; Section 4.11). Source: the analysis program's sensitivity outputs on the frozen coding.

Table E3. Every figure that differs between the analysis on the coding fixed before collection and the final analysis.

Figure	Original analysis (coding fixed before collection)	Final analysis
Training-content summary: related document only (deadline had passed)	6 of 21	4 of 21
Training-content summary: published, identified as the summary (deadline had passed)	7 of 21	11 of 21
Where the summary or related document sits on the provider's site: no document to read (deadline had passed)	7 of 21	6 of 21
Where the summary or related document sits on the provider's site: could not be determined (the crawl was not completed, or not run to a document reached after the window) (deadline had passed)	7 of 21	9 of 21
The document names the model it covers: names specific models (deadline had passed)	12 of 21	13 of 21
The document names the model it covers: no document to read (deadline had passed)	7 of 21	6 of 21
The document states the kinds of training content: no document to read (deadline had passed)	7 of 21	6 of 21
The document states the kinds of training content: lists them (deadline had passed)	10 of 21	12 of 21
The document states the categories of sources: no document to read (deadline had passed)	7 of 21	6 of 21
The document states the categories of sources: lists them (deadline had passed)	11 of 21	13 of 21
The document names particular sources: gives examples of kinds of source only (deadline had passed)	2 of 21	3 of 21
The document names particular sources: names datasets (deadline had passed)	9 of 21	10 of 21
The document names particular sources: no document to read (deadline had passed)	7 of 21	6 of 21
The document refers to the AI Office template: no reference to the template (deadline had passed)	11 of 21	13 of 21
The document refers to the AI Office template: no document to read (deadline had passed)	7 of 21	6 of 21
The document carries a date or a version: dated and versioned (deadline had passed)	10 of 21	11 of 21
The document carries a date or a version: no document to read (deadline had passed)	7 of 21	6 of 21
The document carries a date or a version: neither dated nor versioned (deadline had passed)	1 of 21	2 of 21
Statement on the provider's own web crawling for training: nothing found in the places we looked (all providers)	11 of 26	10 of 26
Statement on the provider's own web crawling for training: web collection stated, crawler not named (all providers)	8 of 26	9 of 26
Training-content summary: summary mentioned, with no location a reader can reach (deadline had passed)	1 of 21	0 of 21
Training-content summary: could not be opened (deadline had passed)	1 of 21	0 of 21
Where the summary or related document sits on the provider's site: could not be opened (deadline had passed)	1 of 21	0 of 21
The document names the model it covers: could not be opened (deadline had passed)	1 of 21	0 of 21
The document states the kinds of training content: could not be opened (deadline had passed)	1 of 21	0 of 21
The document states the categories of sources: could not be opened (deadline had passed)	1 of 21	0 of 21

Figure	Original analysis (coding fixed before collection)	Final analysis
The document names particular sources: could not be opened (deadline had passed)	1 of 21	0 of 21
The document refers to the AI Office template: could not be opened (deadline had passed)	1 of 21	0 of 21
The document carries a date or a version: could not be opened (deadline had passed)	1 of 21	0 of 21
The document names the model it covers: covers the provider in general (deadline had passed)	0 of 21	1 of 21
Served live when read (the documents of Qwen and Tencent on 15 October 2026; Tencent's file opened by the AI-assisted re-check, the identical file saved by the study on 16 October 2026), of those with a document address (deadline had passed)	13 of 14	15 of 15
Not retrieved by our reading, live or from an archive copy, of those with a document address (deadline had passed)	1 of 14	0 of 15
Located documents, deadline had passed: could not be determined (the crawl was not completed, or not run to a document reached after the window)	7 of 13	9 of 15
Located documents, deadline had passed: not reached within three clicks of the home page	2 of 13	2 of 15
Index band: full public record (deadline had passed)	7 of 21	10 of 21
Index band: substantial public record (deadline had passed)	3 of 21	1 of 21
Index band: partial public record (deadline had passed)	1 of 21	2 of 21
Index band: minimal public record (deadline had passed)	8 of 21	6 of 21
Index of Microsoft	77.0 (substantial public record; 0 not determinable)	90.0 (full public record; 0 not determinable)
Index of Qwen	5.0 (minimal public record; 0 not determinable)	57.0 (partial public record; 0 not determinable)
Index of Tencent	0.0 (minimal public record; 7 not determinable)	90.0 (full public record; 0 not determinable)
Index of Thinking Machines Lab	77.0 (substantial public record; 0 not determinable)	90.0 (full public record; 0 not determinable)
Answers not settled, all providers (including answers that could not be opened)	76	70

Note: rows are read by program from the two analyses' frequency, archive, findability, band and index outputs; a value with no row in one analysis counts zero there. The index rows rest on a flagged item (less certain: our two readings agreed below a threshold set in advance on the item this rests on; Section 4.11). Source: the analysis program's outputs on both sets.

Table E4. The correction record: every value the final analysis applies in place of an original one.

Company	Item	Original record	Value the final analysis applies	Under
Cohere	A page about the EU AI Act	Mentioned, with no location a reader can reach	Could not be determined	deviation 33
Google	Statement on the provider's own web crawling for training	Nothing found in the places we looked	Web collection stated, crawler not named	deviation 34
Google	Own reading: marking as far as technically feasible (Article 50(2))	Nothing found in the places we looked	"we are concerned that adding more regulatory complexity — as technical solutions are still evolving — could contradict Europe's goals for competitiveness and simplification"	deviation 34
Microsoft	Training-content summary	Related document only	Published, identified as the summary	deviation 36
Microsoft	The document names the model it covers	Names specific models	Names specific models	deviation 36
Thinking Machines Lab	Training-content summary	Related document only	Published, identified as the summary	deviation 36
Thinking Machines Lab	The document names the model it covers	Names specific models	Names specific models	deviation 36
Tencent	Training-content summary	Could not be opened	Published, identified as the summary	deviation 36
Tencent	The document names the model it covers	Could not be opened	Names specific models	deviation 36
Tencent	Summary or related document served live when read	Not retrieved by our reading, live or from an archive copy	Served live on 15 October 2026 (opened by the AI-assisted re-check; the study saved the identical file on 16 October 2026)	deviation 36
Tencent	The document states the kinds of training content	Could not be opened	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38

Company	Item	Original record	Value the final analysis applies	Under
Tencent	The document states the categories of sources	Could not be opened	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Tencent	The document names particular sources	Could not be opened	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Tencent	The document refers to the AI Office template	Could not be opened	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Tencent	The document carries a date or a version	Could not be opened	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	Training-content summary	Summary mentioned, with no location a reader can reach	Published, identified as the summary	deviation 36
Qwen	The document names the model it covers	No document to read	Names a model family	deviation 36; replaced under deviation 40
Qwen	The document states the kinds of training content	No document to read	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	The document states the categories of sources	No document to read	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	The document names particular sources	No document to read	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	The document refers to the AI Office template	No document to read	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	The document carries a date or a version	No document to read	Could not be determined (the document was not coded)	deviation 36; replaced under deviation 38
Qwen	Summary or related document served live when read	Not asked, because a rule removed the question	Served live on 15 October 2026 (saved by the study's reader)	the note on deviation 36
Tencent	Where the summary or related document sits on the provider's site	Could not be opened	Could not be determined (the link crawl was not run to the document)	the note on deviation 36
Qwen	Where the summary or related document sits on the provider's site	No document to read	Could not be determined (the link crawl was not run to the document)	the note on deviation 36
Tencent	The document states the kinds of training content	Could not be opened	Lists them	deviation 38
Tencent	The document states the categories of sources	Could not be opened	Lists them	deviation 38
Tencent	The document names particular sources	Could not be opened	Names datasets	deviation 38
Tencent	The document refers to the AI Office template	Could not be opened	No reference to the template	deviation 38
Tencent	The document carries a date or a version	Could not be opened	Dated and versioned	deviation 38
Qwen	The document states the kinds of training content	No document to read	Lists them	deviation 38
Qwen	The document states the categories of sources	No document to read	Lists them	deviation 38
Qwen	The document names particular sources	No document to read	Gives examples of kinds of source only	deviation 38
Qwen	The document refers to the AI Office template	No document to read	No reference to the template	deviation 38
Qwen	The document carries a date or a version	No document to read	Neither dated nor versioned	deviation 38

Company	Item	Original record	Value the final analysis applies	Under
Sakana AI	Earliest documented date of outside access to the model	11 September 2026	24 April 2026	deviation 39
Qwen	The document names the model it covers	No document to read	Covers the provider in general	deviation 40

Note: rows are in the order of the correction record; the original record is never edited; the correction record carries both values, and the release's correction-evidence file gives the source of each. Source: the final analysis's corrections file.

AI-assistance disclosure

Collection, the coding passes and drafting for this study are AI-assisted and run under named human editorial control, consistent with our published AI transparency policy (AI Tools Police, n.d.).

- A named person, Mücahit Kaya, is accountable for every published value and every claim in this paper.
- The readings this paper describes, first, second and adjudicating, the AI-assisted re-check after the window, the two content readings of deviation 38, the readings and adjudication of deviation 39, the two readings of deviation 40, and the methods review, are AI-assisted roles under that editorial control. None is a contributor named beside the author.
- Every source cited in the paper's own text, in the register of the Act's text and in the record of the Commission's pages is opened by a person before publication, against a checklist of every reference and its address. Per-value sources in the dataset are not opened one by one by a person; human review covers the record classes the protocol names, and the protocol's checks run over every row.
- Every quotation of the Act comes from the frozen register, checked character for character against a rendered view of the text. The quotation check applies the codebook's straight-quote rule and the whitespace rule for PDF text, including special space characters, and no other.
- The first read of the Commission's pages, on 11 September 2026, was a rendered browser read by the author at all 9 locations. The 2 later rounds were read by the study's automated browser.
- Article 50(4) removes a deployer's disclosure duty for AI-generated text that has undergone human review or editorial control and for which a person holds editorial responsibility. This publication drafts with AI assistance under the named editorial responsibility of Mücahit Kaya and says so on its AI transparency page. Whether that satisfies the paragraph is a legal question this study does not decide for anyone, including itself; no answer turns on it.
- Nothing is written up as first-hand product experience. This study involves no hands-on testing.
- AI assistance never substitutes for the codebook's decision rules.