

[• LEGAL](#)

EU AI Act notice

SumizAI turns a question and an answer into a note, and a language model does the writing. This notice says exactly which parts of the product are AI, who is responsible for what under Regulation (EU) 2024/1689 (the AI Act), and where the output should not be trusted without reading it.

Last updated: 2026-08-14 **Version:** 1.0 **Operator:** 10UP Jan Śliwa, Kminkowa 26, 72-003 Bezrzecze, Poland

01 What this notice is

This is the transparency notice for the artificial intelligence inside SumizAI. It is written for the people who use the product, not for regulators, and it covers Regulation (EU) 2024/1689 — the AI Act — whose transparency rules in Article 50 have applied since 2 August 2026.

The operator of SumizAI is **10UP Jan Śliwa**, Kminkowa 26, 72-003 Bezrzecze, Poland, NIP 8521091009, contact contact@sumizai.com.

This notice is information, not legal advice, and it does not restrict any right you have under the AI Act, the GDPR, Polish law or EU consumer law. Consumers and sole traders buying outside their professional specialisation (*przedsiębiorcy na prawach konsumenta*) keep every statutory right regardless of what any document of ours says.

The word "SumizAI" covers three things and they behave differently, so this notice names them each time it matters: the **website** at sumizai.com, the **desktop application**, and the frozen legacy **web application** at app.sumizai.com. The website contains no AI at all.

02 What the model actually does

SumizAI is built around one loop: you ask, the model answers, and the answer becomes a note. The model is involved at every step of that loop.

It answers your question

The reply you read in the conversation is generated by the model, not retrieved from a database of prepared answers.

It reads part of your vault first

The prompt sent to the provider is not only your typed question. It also carries the table of contents of the vault and the full text of **at most five** notes judged most relevant, within a budget of **24 000 characters**. Which notes were used is shown to you in the application.

It writes the note

The full question-and-answer text is sent again so the model can draft the note: its **title**, its **summary** and its **body** are all model-written. You are not editing a template.

It decides where the note goes

The model places the new note in the table of contents of the vault — the section it belongs under is a model decision, not a rule you configured.

It judges duplicates

The model compares the new material against what is already in the vault and judges whether it duplicates an existing note.

All of that happens automatically, in one pass, without asking you to confirm each step.

Where the call goes. In the **desktop application** the request travels straight from your device to the provider you chose. It does not pass through any SumizAI server, and we never see the question, the answer or the note. In the frozen legacy **web application** at app.sumizai.com the server makes the call for you, decrypting your stored provider key for that one request.

03 Our role: deployer, not model provider

The AI Act distinguishes the party that builds and places an AI model on the market from the party that puts it to use. We are the second one.

- We are a **deployer**: we integrate general-purpose models built by other companies into a note-taking application.
- We are **not a provider of a general-purpose AI model**. We do not train, fine-tune or host a model. SumizAI ships with no AI capability of its own.
- The model is **your choice**, made in the settings, among seven options: Anthropic, OpenAI, Google Gemini, Groq, OpenRouter, a local Ollama instance, or any other OpenAI-compatible endpoint you enter yourself.
- It runs on **your own API key**. You bring the key; the application has no account, quota or credit with any provider. A local Ollama may need no key at all — and on the desktop, where the call starts on your own machine, a local Ollama means nothing leaves it.

The company whose model you selected is the provider of that model and carries the provider duties for it, including the duties for general-purpose AI models. Your relationship with that company is your own: their terms, their pricing, their data handling. We are not a party to it, and in

the desktop product we could not inspect that traffic even if we wanted to.

04 Risk classification: limited risk

SumizAI is a **limited-risk** AI system. The obligation that applies to it is transparency under Article 50 — telling you that you are dealing with AI and that the notes are AI-written. That is what this document does.

It is not a high-risk system under Annex III. Annex III lists the areas where an AI system decides something consequential about a person. SumizAI does none of them:

- It does not perform biometric identification, biometric categorisation or emotion recognition.
- It is not used to manage or operate critical infrastructure.
- In education it does not decide admission, does not evaluate learning outcomes, does not assign a level, and does not monitor anyone during an exam. It writes notes for the person who typed the question, and nobody is graded by it.
- It plays no part in recruitment, task allocation, promotion or termination.
- It gives no access to any public benefit or private service, scores no creditworthiness, prices no insurance and triages no emergency call.
- It is used in no law-enforcement, migration, border or asylum context.
- It has no role in the administration of justice or in democratic processes.

It is also not a safety component of a regulated product under Annex I. Because it is not high-risk, there is no conformity assessment, no CE marking, no EU database registration and no notified body involved — and we claim none of those. We hold no certification of any kind for this product.

05 Transparency obligations and how we meet them

The table below lists each transparency point, what its status is here, and how it is met. Where something is not implemented, it says so.

OBLIGATION	STATUS	HOW IT IS MET
Article 50(1) — people must know they are interacting with an AI system	Met	The product is presented as an AI note-taker throughout: you pick the provider and the model yourself, you paste your own key, and the application shows which notes were fed to the model for each answer.

OBLIGATION	STATUS	HOW IT IS MET
Article 50(2) — machine-readable marking of synthetic content	Not implemented by us	Notes are saved as ordinary Markdown files, and we do not add a watermark or any machine-readable marker identifying them as AI-generated. Whether the model provider marks its own output is the provider's matter, not something we can promise. Treat every note in a SumizAI vault as AI-written.
Article 50(3) — emotion recognition and biometric categorisation	Not applicable	Neither is performed. The product works on text you typed and notes you already have.
Article 50(4), first subparagraph — deepfake disclosure (a deployer duty)	Not applicable	The product generates text only. There is no image, audio or video generation anywhere in it.
Article 50(4), second subparagraph — disclosing AI-generated text published to inform the public on matters of public interest (a deployer duty)	Not applicable	SumizAI publishes nothing. A note is written into your own vault: on the desktop into a file on your own disk, and in the frozen legacy web application into the storage of your account there. Nothing is sent to any audience.
Article 50(5) — the information must be clear and given at the latest at the first interaction	Met by publication	This notice is published on sumizai.com and linked from the footer of every page, next to the privacy policy and the terms of service. That is what we can stand behind; we do not claim that the application itself repeats this notice when you first open it.
Article 4 — AI literacy of the people using the system	Addressed in writing	This notice explains what the model does, what is sent to the provider, and where the output fails. We run no training programme beyond documentation.
Conformity assessment, CE marking, registration, certification	None — and none claimed	These apply to high-risk systems. SumizAI is not one, so no such procedure was carried out and no such claim is made anywhere in our documents.

06 Human oversight: you are the reviewer

Nothing the model writes is published or shown to an audience. A note is generated and saved into your own vault — on the desktop into files on your own disk, in the frozen legacy web application into the storage of your account there — and there it waits for you.

- You read the answer before it becomes anything.

- You can edit the note, rewrite it, or delete it. On the desktop the notes are plain Markdown files in a folder you chose, so any editor works on them.
- You can move a note if the model filed it in the wrong place, and keep a note the model called a duplicate.
- The only outbound request the AI makes is the call to the provider you chose, carrying the material described above. SumizAI posts nothing, publishes nothing and shares nothing with anyone else on your behalf.

Be honest with yourself about one thing: **the product does not force you to review anything.** There is no mandatory approval step today — the note is written and stored in a single pass. The oversight is real but it is yours to exercise, and it matters most when you later rely on a note you never actually read.

No decision is taken about you. The model writes text; it does not score, profile, rank or evaluate any person, and nothing in SumizAI produces a decision based solely on automated processing that has a legal or similarly significant effect on you within the meaning of Article 22 GDPR. Which personal data exists, who is its controller and how you exercise your rights of access, rectification, erasure, restriction, objection and portability is set out in the [Privacy Policy](#). If your notes or questions contain personal data of other people, sending them to your chosen provider is your decision and your legal basis.

07 Model choice, versions, and your provider's terms

There is no version pinning. You choose a provider and a model name; the company behind that name can change the model it points at, retire it, or alter its behaviour at any time, without telling us. The same question can therefore produce a different answer tomorrow, and a note written last month was written by whatever that model was last month. We cannot reproduce a past output and do not promise stability of results.

Whether your content is used for training is decided by your provider, not by us. The prompt — your question, the table of contents, and the notes selected as context — is sent under the account whose API key you pasted. What that company does with it is governed by the terms of that account. Read them. Many providers treat paid API traffic differently from their consumer chat products, and some offer a setting for it; the setting lives with them, not in SumizAI.

Practical consequences:

- If you do not want your notes to leave your machine at all, use the desktop application with a local Ollama model: the call is made by your own device and no request to a third party happens. In the frozen legacy web application the notes are already stored on a server and the call is made there, so a local model on your own computer is not what happens.
- If you point the custom option at your own OpenAI-compatible endpoint, that endpoint is yours to assess — we know nothing about it.

- Changing provider changes the answer style, the duplicate judgement, and the table-of-contents placement. They are all model decisions.

08 Prohibited practices (Article 5): none are performed

Article 5 of the AI Act bans a short list of practices outright. SumizAI performs none of them, and we confirm that plainly:

- No subliminal, manipulative or deceptive techniques that distort behaviour.
- No exploitation of age, disability or a social or economic situation.
- No social scoring of people or groups.
- No prediction of criminal offences based on profiling or personality traits.
- No untargeted scraping of facial images to build a recognition database.
- No emotion inference in the workplace or in education.
- No biometric categorisation to infer race, political opinions, trade-union membership, religion, sex life or sexual orientation.
- No real-time remote biometric identification in publicly accessible spaces.

The product accepts typed text and reads notes from a vault you control. It has no biometric input, builds no profile of you, and takes no decision about any person.

09 What the AI gets wrong

Language models produce fluent text whether or not the text is true. This is the section to remember.

- **Hallucination.** The model can invent facts, dates, sources, quotations and citations, and it will present them in the same confident tone as everything else. A SumizAI note is not a verified document.
- **A limited view of your vault.** Each answer is built from the table of contents plus at most **five** notes, capped at **24 000 characters**. If the answer depended on a sixth note, or on the part of a long note that did not fit inside the budget, the model never saw it. The application shows you which notes it used — that list is worth checking.
- **Wrong duplicate judgements.** The model can call a genuinely new note a duplicate, or fail to notice a real repetition.
- **Wrong placement.** A note can land under the wrong heading in the table of contents.
- **No accuracy guarantee.** We make no promise, express or implied, that any answer, title, summary or note body is correct, complete or fit for a particular purpose.

- **Not professional advice.** Do not use SumizAI output as the sole basis for a medical, legal, financial, tax, safety or other consequential decision. Verify anything that matters against a source you trust.

The statutory rights of consumers and of *przedsiębiorcy na prawach konsumenta* under Polish and EU law are not affected by this section.

10 How this fits with our other documents

This notice explains the AI. The neighbouring documents explain the rest:

- [Privacy Policy](#) — what data exists, where it is, and who is the controller of it. The short version: the desktop application keeps everything on your device and sends us nothing; the website uses analytics after consent; the legacy web application at app.sumizai.com does store accounts and notes on a server.
- [Terms of Service](#) — the contract, what you may do with the product, and the limits of our liability.
- [DPA](#) — the data processing agreement, which is only meaningful for the web application, because in the desktop product we process nothing on your behalf.
- [Cookie policy](#) — the website's analytics and the strictly necessary cookies of the web application. The desktop application sets no cookies.
- [Refunds](#) — the published offer and your withdrawal rights.

If this notice and another document appear to contradict each other on a point about the AI, tell us at contact@sumizai.com and we will fix the wording rather than argue about which one wins.

11 Review and changes

The facts in this notice were verified against the product source on **14 August 2026**, and the version and date are printed at the top of the page.

We review it when the product changes in a way this document describes — a new provider, a change to how much of the vault is sent, a change to what the model decides on its own — and when the AI Act or its guidance changes what has to be said. Superseded statements are replaced, not quietly softened: if a duty moves from "not implemented" to "met", the table will say so.

Continuing to use SumizAI after a change means you use the version then published. Nothing in this notice creates an obligation on you to accept anything; it is a description of the system.

12 Contact

Questions about the AI in SumizAI — what is sent, what the model decides, why a note came out the way it did — go to a human:

- **10UP Jan Śliwa**, sole proprietorship registered in Poland (CEIDG)
- Kminkowa 26, 72-003 Bezrzecze, Poland
- NIP 8521091009
- contact@sumizai.com

Questions about the model itself — its training data, its safeguards, its own AI Act documentation — belong to the company whose model you selected. We can tell you what SumizAI sends and when; we cannot answer for what happens inside somebody else's model.

Two complaint routes exist and they are not the same one. If your complaint concerns **personal data** — what is stored, what was sent, a request of yours we failed to answer — you may lodge it with the Polish supervisory authority, the President of the Personal Data Protection Office (Prezes Urzędu Ochrony Danych Osobowych), ul. Stawki 2, 00-193 Warsaw. If your complaint concerns the **AI system itself**, the AI Act gives you the right to complain to the competent national market surveillance authority; in Poland that is the authority designated by the national law implementing the Regulation. Neither route costs you the right to write to us first, and neither is limited by anything in this notice.

We publish no telephone number, so e-mail is the channel.