Artificial intelligence

Version of November 19, 2026

What WYRLO does with artificial intelligence, what it does not guarantee, and what you must tell your own users. Information page under Regulation (EU) 2024/1689 on artificial intelligence; the contractual commitments are set out in the terms of use.

Contents

- What WYRLO does
- Models used
- What WYRLO does not guarantee
- European AI Regulation
- What you must do with regard to your users
- Questions and reports

What WYRLO does

- Verification: for each answer from your RAG, WYRLO detects prompt injections with rules, compares the answer with the passages provided and, in French and English, has a Mistral AI language model check that every statement appears in the passages. From this it derives a confidence level and a decision: answer, or say “I don’t know”.
- Zero-copy index: a Mistral AI embedding model turns the passages into numerical representations, for search. The text is not kept.
- Drafting (/v1/answer, in preview): a language model drafts a cited answer from your passages, then WYRLO verifies it. The model is the one managed by WYRLO (Mistral AI) or the one you choose.

- Nightly improvement: from each project's traces, WYRLO proposes and tests settings (vocabulary, passages kept, confidence thresholds). A change is only adopted if it improves the measured decisions without worsening any; it is then monitored and reverted automatically if it worsens the result. A language model drafts the explanation of failures and the report.
- Help assistant: in the developer space, a language model answers your questions about WYRLO based on the help centre, and its answer is verified before being displayed.

Models used

- Mistral AI (France), through its European endpoint: language models (mistral-small, mistral-large) and embedding model (mistral-embed). They only receive texts masked by WYRLO.
- The model of your choice, for /v1/answer: OpenAI, Anthropic, Mistral AI, Azure OpenAI or a compatible address, called with your key, without masking.
- Rules written by WYRLO (injections, masking, consistency checks), which decide on their own when a model is not available.
- WYRLO trains no model with your data, neither its own nor a third party's.

What WYRLO does not guarantee

- The accuracy of an answer: an answer marked "certain" may be wrong, and an "I don't know" may set aside a good answer. The confidence level is an estimate.
- Verification only judges the agreement between the answer and the passages provided: if your documents are wrong, outdated or incomplete, an answer faithful to these documents will be too.

- Verification by a language model only exists in French and English; in other languages, the confidence level relies on other, less reliable signals.
- When WYRLO does not respond in time, the plan's volume is reached or a model is unavailable, the answer goes through marked "unverified".
- Injection detection and masking are automatic and not exhaustive; people's names are not masked.

European AI Regulation

WYRLO provides an AI service to professionals, who integrate it into their own applications. It uses general-purpose AI models from Mistral AI, which is their provider. WYRLO is not designed for the high-risk uses listed in Annex III of the Regulation (employment, education, credit, essential public services, law enforcement, justice, migration, etc.) and takes no decisions about people: it only decides whether an answer can be shown.

The transparency obligations of article 50 of the Regulation towards people who interact with an AI, or read a text generated by an AI, rest first with the organisation that makes the application available to them, that is, you.

What you must do with regard to your users

- Inform them, clearly and at the latest at their first interaction, that they are interacting with an AI system, unless this is obvious; and that an answer drafted by /v1/answer is generated by an AI.
- Explain to them what "I don't know" and "unverified" mean, and offer them recourse to a human when the stakes justify it.

- Not use WYRLO as the sole basis for a decision producing legal effects on a person, nor in a system prohibited by article 5 of the Regulation; inform us before any integration into a high-risk system.
- Train the people in your organisation who use or integrate WYRLO on its limits (AI literacy, article 4 of the Regulation).
- Inform your users of the processing of their data (questions, answers) and of the use of WYRLO and its subprocessors, and not send sensitive data without masking it.
- If you are a Partner, pass these obligations on to your End clients.

Questions and reports

For any question about AI in WYRLO, or to report a problematic answer or a risk, write to contact@wyrlo.io.

Terms of use Security and data Subprocessors