

AI Transparency Notice — UnderStack Code

Version 1.0 · 25 September 2026

This notice explains how UnderStack Code uses artificial intelligence, as required by article 50 of the EU Artificial Intelligence Act (Regulation (EU) 2024/1689).

You are interacting with an AI system

UnderStack Code is an AI system. Its answers, plans, code, reviews, commands and images are produced by AI models, not by a person. The app tells you this on first launch and in Settings > About & legal.

Which models

The models run locally on your Mac and are downloaded from their publishers on your request:

Use	Default model	Publisher	Licence
Chat, planning, code, review	gpt-oss-20b (through Ollama)	OpenAI	Apache 2.0 + gpt-oss usage policy
Reading screenshots and photos	Gemma 3 12B (through Ollama)	Google	Gemma Terms of Use
Creating images	Z-Image Turbo, 4-bit (through MFLUX)	Tongyi-MAI (Alibaba)	Apache 2.0

You can choose other local models in Settings > Models. UnderStack does not train these models and does not use your data to train anything.

AI-generated images are marked

Every image created with ✦ Image is marked as AI-generated, in a machine-readable way, inside the PNG file:

- XMP metadata with the IPTC digital source type `trainedAlgorithmicMedia`, the generator (UnderStack Code) and the model;
- PNG text fields (`Software` , `Description` , `Comment`) stating that the image is AI-generated with UnderStack Code and naming the model.

The chat also labels each image as AI-generated. The marking is kept when you save the image to your project. Some tools remove metadata when they edit or re-encode images; if you publish an image that could be mistaken for a real person, place or event, say that it is AI-generated.

Limitations

- AI output can be wrong, outdated, insecure or biased, and can resemble existing works.
- The app does not know whether generated code is correct until it has been tested; it runs your project's tests and an independent AI review, but these do not replace your own review.
- Image generation can produce errors in text, hands, faces and details, and must not be used to deceive.

Human oversight

The model proposes; you decide. The app classifies every action by risk and asks for your approval before editing files, running commands, installing packages or publishing previews. Privileged, production and destructive actions always need explicit confirmation, and some actions are blocked entirely. Every decision is logged under Settings > Activity.

Your data

Inference happens on your Mac. Prompts, files and images are not sent to UnderStack or to the model publishers. See the [Privacy Policy](#).

Contact

info@understack.dk · UnderStack, Aarhus, Denmark · CVR 46327608