AI Model Summary

Last modified: 07 May 2026

1. Scope and Purpose

This document ("Summary") is issued by:

Company Name	Tengrai Artificial Intelligence Korlátolt Felelősségű Társaság
Headquarters	Hungary, 6724 Szeged, Rókusi-boulevard 21, 1st floor, door 4
Company Reg. No.	06-09-028918
Tax number	32315028-2-06

("Provider")

for the public and the use of competent supervisory authorities within the meaning of Regulation (EU) 2024/1689 on Artificial Intelligence (the "AI Act") and any other applicable statutory instrument governing the placing on the market, or commissioning, of general-purpose artificial intelligence systems. The Provider's Proprietary Models are capable of generating images. Third-Party Models accessible through the Platform may additionally generate video and text.

Its purpose is to set out, in a legally cognisable manner, (i) the technical characteristics of the Provider's general-purpose multimodal generation models (the "System"), (ii) the data governance measures applied during training, and (iii) the ex-ante and ex-post risk-management and redress mechanisms that ensure the System's continuous conformity with European Union and Member-State law.

The Provider's GPAI obligations are governed by Chapter V, Articles 51–53 of the AI Act, which entered into application on 2 August 2025. The Provider is actively evaluating adherence to the GPAI Code of Practice published by the European Commission in July 2025 as a voluntary compliance instrument under Article 53.

2. Definitions

For the avoidance of doubt, the following capitalised terms shall have the meanings assigned to them below. All other terms shall bear the meanings attributed in the AI Act unless the context otherwise requires.

Hyperalign Architecture the hidden alignment, constraint-enforcement and policy-verification subsystem embedded in the System, described in Section 5.

Policy Store the machine-readable repository of binding legal, ethical and contractual rules consulted by Hyperalign.

Public-Source Dataset any image-text dataset released under a licence permitting commercial use, modification, and redistribution without royalty, including but not limited to those enumerated in Section 4.

Disallowed Material any visual or textual content the generation, possession or dissemination of which would infringe Union or Member-State law or the Provider's own acceptable-use policies.

Proprietary Models the Provider's own text-to-image generation models developed, trained, and operated by the Provider, accessible directly on the Platform.

Third-Party Models general-purpose AI models developed and owned by third-party providers, made available to users through the Platform as a frontend interface.

3. High-Level System Description

3.1 General Function

The System encompasses two categories of models:

a) Proprietary Models

The Provider's own diffusion and transformer-based text-to-image models, which generate images from natural-language text prompts. These models are developed, trained, and operated by the Provider, and are accessible directly on the https://alberti.studio and https://albertistudio.ai web applications.

b) Third-Party Models

General-purpose AI models developed by third-party providers, made available to users through the Platform as a frontend interface. For these models, the Provider acts as a downstream deployer. Data governance, training provenance, and model-level technical documentation for Third-Party Models are the responsibility of their respective upstream providers.

3.2 Functional Blocks (Proprietary Models)

1. **Prompt Reception (S1) —** natural-language text prompts are accepted; no pre-filter is applied at ingestion.
2. **Primary Generation (S2) —** a base diffusion model constructs a provisional image (O_raw).
3. **Hidden Constraint Enforcement (S3) —** Hyperalign evaluates O_raw against the Policy Store and, where necessary, deletes, replaces or regenerates non-compliant segments.
4. **Delivery (S4) —** the transformed image (O_final) is transmitted to the user; no refusal text, apology or censorship notice is ever surfaced.

3.3 Third-Party Model Infrastructure

The Provider makes Third-Party Models available to users through Replicate, Inc., a third-party model infrastructure provider. Models made available through Replicate include, among others, models under the NanoBanana family and additional models that may be added over time. Users of the Platform may generate images, video, and text through these models subject to both the Provider's own acceptable-use policies and the upstream terms applicable to each model.

Relevant Replicate legal documents:

Terms of Service	https://replicate.com/terms
Privacy Policy	https://replicate.com/privacy
Acceptable Use Policy	https://replicate.com/acceptable-use-policy

The Provider does not control the training data, architecture, or model weights of Third-Party Models. Compliance obligations relating to training data provenance, copyright, and model-level technical documentation for those models are the responsibility of their respective upstream providers. The Provider's obligations with respect to Third-Party Models are those of a deployer under the AI Act, including ensuring that use of such models through the Platform remains within the bounds of the Provider's acceptable-use policies and the AI Act's prohibitions.

4. Training-Data Provenance and Licensing (Proprietary Models only)

4.1 Source Datasets

The Provider certifies that the image corpus utilised for pre-training and fine-tuning of its Proprietary Models exclusively comprises Public-Source Datasets and licensed Private-Source Datasets (collectively, the "Training Set").

The Provider attests that the Training Set consists exclusively of image–text corpora that are lawfully and openly distributed under permissive copyright instruments, namely Creative Commons Attribution licences (CC BY 4.0, CC BY 2.0), public-domain waivers (CC0, Public-Domain Mark), or equivalent irrevocable custom licences granting worldwide, royalty-free rights of use, modification and commercial redistribution. Representative sources include, inter alia: the MS COCO reference benchmark (CC BY 4.0); Google's Open Images collection (images predominantly CC BY 2.0); the Conceptual Captions and Conceptual 12M web-alt-text corpora (images ≥ CC BY 2.0); the 99-million-item YFCC-100M multimedia archive released under a spectrum of Creative Commons variants; and the Wikimedia Commons, Smithsonian Open Access and other public-domain repositories whose content is either CC-licensed or free of re-use restrictions.

Prior to ingestion, the Provider conducted automated licence-tag verification and eliminated every record bearing a Non-Commercial, No-Derivatives or otherwise incompatible encumbrance, thereby ensuring that every retained image satisfies the AI Act Article 53(1)(c) requirement for copyright compliance and the obligation to publish a sufficiently detailed summary of training content under Article 53(1)(d). Comprehensive download manifests, cryptographic checksums and licence details are preserved in the Dataset Register and may be made available to the competent authorities upon reasoned request.

4.2 Exclusion of Proprietary Content

No images obtained from closed social-media platforms, privately hosted repositories, or otherwise encumbered sources were ingested. No scraping tools bypassing technical protection measures were employed.

5. Data Due-Diligence, Cleansing and Risk Mitigation (Proprietary Models only)

Given that the Training Set exceeded seven (7) billion candidate images, manual per-image review would have required in excess of 300 years at one (1) second per image; such review was therefore manifestly impracticable. The Provider accordingly adopted the multi-layer automated procedure set out below:

1. Metadata Pre-Screening — textual metadata (captions, alt-text, URLs) were scanned for lexical indicators of obscenity, extremist ideology or the sexual exploitation of minors; flagged entries were discarded at source.

2. **NSFW Visual-Classifier Pass —** a ResNet-based classifier, fine-tuned on the "Journey NSFW" and "Open NSFW" benchmarks, automatically flagged any image scoring above a 0.25 probability of sexual content or graphic violence.

3. **Perceptual-Similarity Filtering —** CLIP embeddings were k-nearest-neighbour-compared against a block-listing gallery of disallowed exemplars (e.g. known extremist insignia, CSA material); hits were expunged.

4. **Semantic-Similarity Filtering —** captions were embedded via SBERT; cosine similarity to a curated vector bank of illicit semantics triggered exclusion.

5. **Deduplication and Quality Control —** near-duplicates (cosine similarity > 0.999) and corrupted or blank files were removed.

6. **Iterative QA Team Refinement —** periodic audits located residual infractions; the corresponding heuristic or model weights were re-trained and the pipeline rerun.

The foregoing measures resulted in the elimination of millions of images, approximately 3.8% of total candidate images, including all entries flagged as high-risk via LAION's own tags.

6. Hyperalign Constraint-Enforcement Mechanism (Proprietary Models only)

6.1 Legal and Technical Basis

Hyperalign constitutes an alignment module that supervises image generation in real time by reference to the Policy Store. Enforcement is performed token-by-token for diffusion latents and, where required, by post-generation in-painting. The architecture thereby:

- accepts every text prompt without exception
- suppresses any observable refusal signal
- delivers only compliant image output, thereby collapsing the success probability of gradient-free jailbreak attempts to the failure probability of the alignment classifier.

6.2 Extensibility and Human Oversight

The Policy Store is maintained in YAML format, version-controlled, and amendable by compliance personnel without necessitating retraining of the core model. Consequently, amendments to the AI Act, forthcoming delegated acts, or Member-State-specific prohibitions can be incorporated within a maximum of ten (10) working days, after which the revised rules bind all subsequent generations.

7. Ongoing Governance, Monitoring and Accountability

1. **Logging and Traceability —** a tamper-evident audit log records each text prompt, the SHA-256 hash of O_final, and the precise policy clauses invoked during any alteration.

2. **Periodic QA Exercises —** at least semi-annually the Provider conducts adversarial testing against emergent jailbreak vectors (e.g., multilingual obfuscation).

3. **Incident Response —** suspected policy-enforcement failures trigger an immediate suspension of the affected endpoint and a root-cause analysis within forty-eight (48) hours.

4. **User Feedback Channel —** any visitor may file a notice of alleged unlawful output via https://alberti.studio; substantiated claims compel retraining or policy-file amendment.

8. Conformity with the AI Act

The Provider's Proprietary Models constitute general-purpose AI models within the meaning of Article 3(63) of the AI Act. The following sets out the Provider's conformity position against the applicable GPAI obligations under Chapter V.

Article 51 (Classification) the Provider has assessed the cumulative compute used in training the Proprietary Models. The training threshold for systemic risk (10^{25} FLOPs) has not been reached; accordingly, the additional obligations of Article 55 do not apply at this time. The Provider monitors this position as model capability evolves.

Article 53(1)(a) (Technical Documentation) this Summary, together with detailed internal design dossiers and the Dataset Register, constitutes the Provider's technical documentation made available to regulators upon reasoned request.

Article 53(1)(b) (Downstream Information) the Provider maintains information and documentation enabling downstream integrators to understand the capabilities, limitations, and acceptable-use constraints of the Proprietary Models.

Article 53(1)(c) (Copyright Compliance Policy) the Provider maintains a copyright compliance policy governing training data selection, as summarised in Section 4, ensuring conformity with Directive 2019/790/EU on copyright in the Digital Single Market.

Article 53(1)(d) (Training Data Summary) Sections 4 and 5 of this document constitute the Provider's public summary of training content, including sources, licensing instruments, and cleansing methodology.

Article 50(4) (AI-Generated Content Transparency) outputs generated by the Proprietary Models are subject to machine-readable marking in conformity with the Provider's content-provenance obligations. Users are informed through the Platform's terms and conditions that content is AI-generated.

GPAI Code of Practice the Provider is actively evaluating formal adherence to the GPAI Code of Practice (July 2025) as the primary voluntary compliance instrument under Article 53. Until formal adherence is confirmed, the Provider demonstrates compliance through the alternative means set out in this Summary.

Guidelines for Secure AI System Development (UK NCSC / CISA) the System's silent, policy-aware output filtering accords with the non-interactive defence pattern endorsed therein.

9. Limitations and Disclaimer

Nothing herein shall be construed as a warranty of absolute censorship or error-free performance. While Hyperalign materially mitigates unlawful output, residual risk cannot be excluded. The Provider shall, however, employ best endeavours to rectify any substantiated deficiency forthwith upon notice.

The technical documentation in this Summary applies exclusively to the Provider's Proprietary Models. For Third-Party Models made available through the Platform, users and regulators are directed to the upstream providers' own technical documentation and compliance disclosures.

10. Contact Details

Website: https://alberti.studio / https://albertistudio.ai
E-mail: support@alberti.studio
Postal address: Hungary, 6724 Szeged, Rókusi-boulevard 21, 1st floor, door 4