DOTCHECK

Technical Report

Architecture, pipelines, and measured holdouts

VERSION	Docs v2026.10
DATE	3 September 2026
ENGINES	Vermeer · Valla · Muybridge · Helmholtz

The figures in this document come from held-out test sets. They estimate how AI-like content looks under our models — they are not proof of authenticity.

Petr Jaroch (trading as DotCheck), IČO 03330389, Jenikovice 2, 535 01 Prelouc, Czech Republic

dotcheck.ai/docs

Contents

1. Scope

This report specifies the live in-house scoring path: Purkyne Gate, the named engines (Vermeer, Valla, Muybridge, Helmholtz), and the holdouts those engines passed. Use it to cite the gates or to wire against the HTTP contract.

It is not a training recipe and not a legal authenticity opinion. The companion **Model Card** on the same docs page lists the same holdout numbers in short form. Same engines on Free, Premium, and Pro.

2. Score

A DotCheck score is a probability in $[0, 1]$. Higher means the content looks more AI-generated under the product path for that media type. Soft labels such as “likely human” are a reading of that number.

The product UI shows **one** product percent. Purkyne Gate may lock that number at 1.0 or 0.0 before a model runs. Reports, Dossiers, Verify, and the issued PDF may name the basis **Purkyne Gate** when an AI-origin lock applied.

The number is an estimate on the kinds of content we tested. It is not a certificate of authorship. A 0% capture lock is not a claim that the photographed scene is true. Do not use a score alone for high-stakes automated decisions about a person.

3. System

Clients talk to Express. Express owns the product score: content-hash cache, Purkyne Gate on **raw** media, Covenant on video-with-sound, and optional Pro confirm gating. The inference service runs named models only when the gate stays open. It does not apply Purkyne and it does not assemble Covenant.

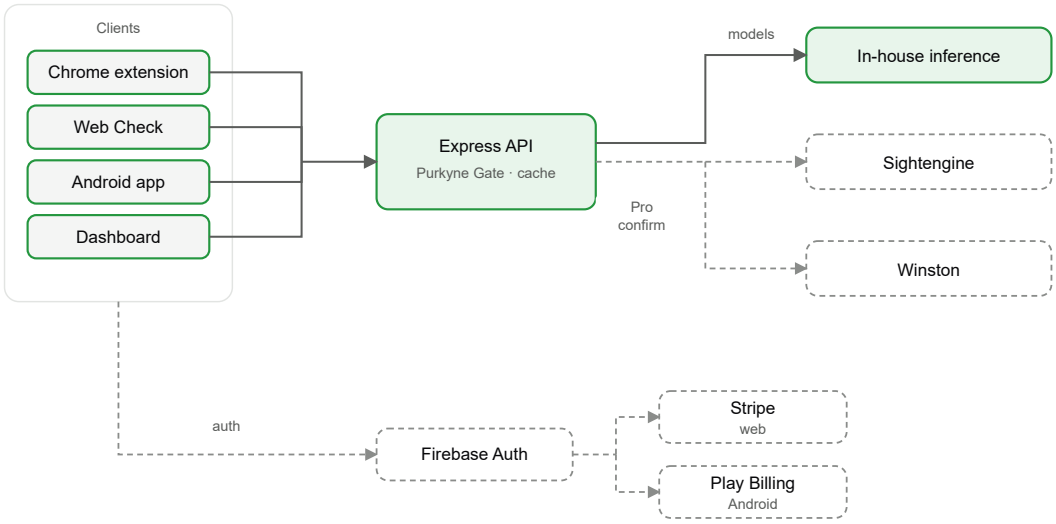

Figure 1 — Runtime: clients ⟶ Express (Gate, cache, Covenant) ⟶ in-house inference. Confirm vendors are explicit-only.

Models run on DotCheck-operated hosts. Live heads are also published under Apache-2.0 on Hugging Face. If inference is unavailable, the API returns HTTP 503. It does not silently send Free or Premium traffic to a paid vendor.

PACKAGE	OWNS
Chrome extension	Page scan, studio/place context; scores via Express
Website	Check, Dashboard, docs; raw upload for Purkyne
Android app	OS Share and on-device Check; raw upload; same Express path
Express API	Product score: Gate, cache, Covenant, confirm gating, reports
Inference	FastAPI scorers; npz heads; CPU serve; no Purkyne; 503 if a required head is missing
Cloud Functions	Stripe and Play subscription status

3.1. HTTP surfaces

ROUTE	ROLE
`POST /analyze-pair`	Still pair: 256 transport + 224 center; Vermeer@14.2 after Gate
`POST /analyze-image`	URL stills and Confirm; pair-encoded on the server; GIF ⟶ video
`POST /analyze-text`	Eight languages; Winston confirm on Pro ask
`POST /analyze-video` / frames	pair2 stamps; Muybridge@4.4 stamp-cap max; GIF ≡ video
`POST /analyze-audio`	Helmholtz@3.1 — file on every plan; live tab Premium/Pro
`POST /analyze-document`	PDF/DOCX; text + capped images; Gate on embeds
`POST /upload-and-analyze(-video)`	Check/Android raw multipart for Purkyne
`POST /share` + status/download	Clip from a real product score already in cache
`POST /report` · `GET /report/:id` · `POST /report/verify`	Issue from a server-held score; Verify freezes basis
`GET /usage`	Remaining analyses and Pro confirms

4. Product path

A request is scored once. The client sends evidence; Express assigns the probability. The client never supplies the number.

On a cache hit, Express returns the stored product score. On a miss, Purkyne Gate runs on raw bytes (and any honest place context). If the gate stays open, named models score. When a clip has usable sound, Helmholtz sits beside Muybridge (**Covenant**; no public fused claim table). Confirm runs only if a Pro user asks after a product score exists.

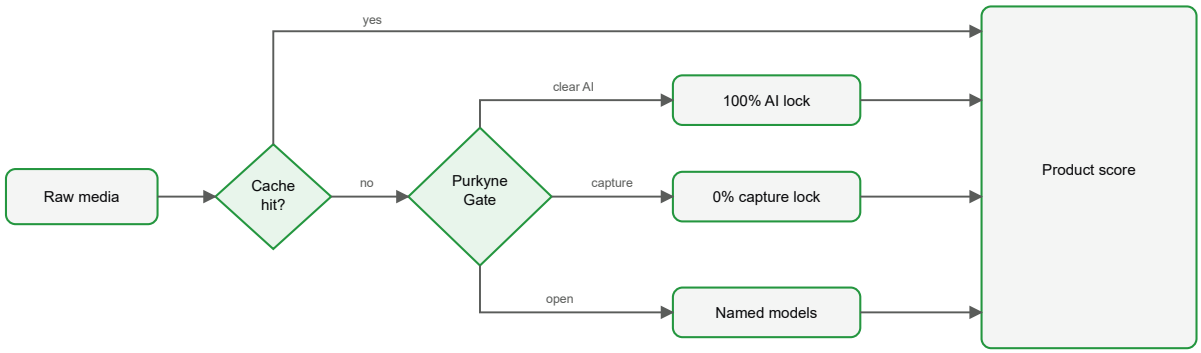

Figure 2 — Cache, then Purkyne Gate on raw media, then named models if the gate stays open.

4.1. Purkyne Gate

Express owns the gate. Inference never applies it. Input is **raw** bytes (original file and, on Check/Android, original filename). A max-side-256 JPEG prepared for the models strips provenance — it is never Gate input. The Chrome extension may send honest place context (allowlisted AI-studio tiles, platform AI disclosure).

OUTCOME	WHEN	PRODUCT SCORE
Place	Allowlisted AI-studio result, or a platform AI disclosure next to the media	100% · lasting
File origin	High-confidence generator credentials / C2PA / AIGC / related origin marks	100% · lasting
Capture	Native camera or phone print still present on an original upload	0% · lasting

If origin and capture both fire, 100% wins. Weak or spoofable marks fail closed: the gate stays open and a named model scores. A Gate lock is stored under the content hash and re-shown later. Confirm never writes a Gate stamp.

4.2. Scoring constraints

CONSTRAINT	RULE
GIF	Video path. Never a still first-frame shortcut.
Confirm	After a product score, on ask. Not the Pro API. Never video/GIF. Never a Gate stamp. 300 / UTC month.
Language	`unsupported_language` — not scored.
File audio	Every plan.
Live-tab audio	Premium and Pro.
Inference down	HTTP 503. No silent vendor. No silent frames-only if soundtrack was sent and failed.

5. Pipelines

The subsections below start after the gate stays open.

5.1. Image

Two views from the **original** pixels:

VIEW	GEOMETRY	JPEG
Transport	Max-side 256, aspect kept, no pad	q70
Center	Native 224×224 center crop	q85

Short side under 224 is refused. Cache key: `sha256(transport + "|" + center)` with engine `inhouse@14`. Inference embeds transport on the 13.12 SigLIP spine with Niepce 13.14 and center on the 6.6 spine with Janssen 6.6; the 14.2 stack returns one P(AI). Serve: `/v1/analyze-pair` `kind=still`.

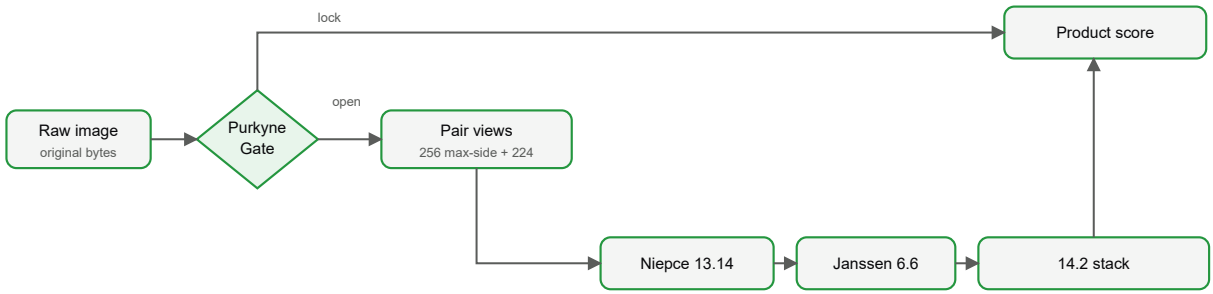

Figure 3 — Image: pair views, Niepce 13.14 + Janssen 6.6, 14.2 stack.

5.2. Text

Supported: English, Chinese, Spanish, Portuguese, French, Italian, German, Dutch. Express detects language first. Other input, and Traditional-Chinese-dominant text, is `unsupported_language`.

English, Spanish, Portuguese, French, Italian, German, and Dutch share Rask and Comenius plus a named tiny; a language-specific stack maps those features to P(AI). Chinese is Ma Jianzhong — independent of that trio. Simplified Chinese only.

Plain text has no container print; the language gate is the first hop. Studio **place** can still lock 100% when the extension scores result content on an allowlisted AI studio.

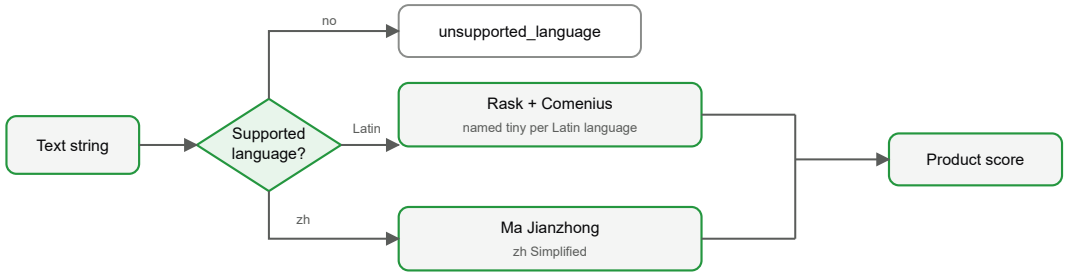

Figure 4 — Text: language gate; Latin heads share Rask + Comenius + tiny; Chinese is Ma Jianzhong.

5.3. Video

Stamps, not full temporal forensics. GIF is this path.

pair2 samples two inner frames on clips up to three minutes, then more on longer clips (cap 12). Each stamp uses the same pair geometry as stills. Muybridge@4.4 runs Plateau 3.8 on transport and Blau 1.4 on center, then the 4.4 stack. Clip fold is stamp-cap (max). If the pair stack is missing, scoring stops with a clear error — still heads are not reused. A usable soundtrack can add Helmholtz windows (Covenant); silence stays frames-only Muybridge.

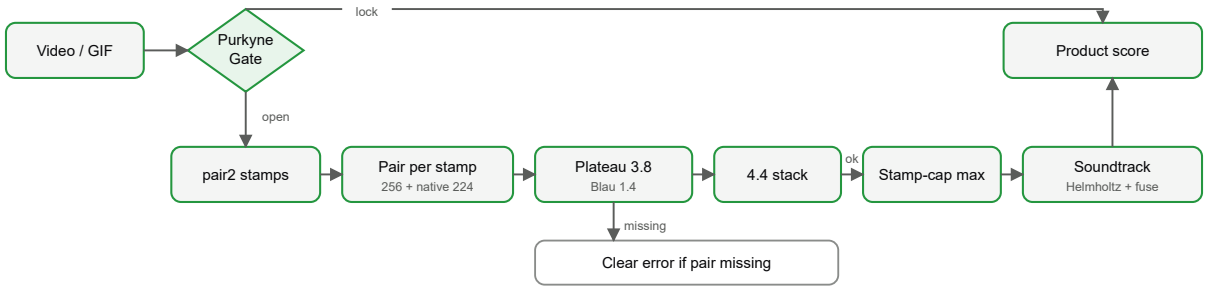

Figure 5 — Video: pair2 stamps, 4.4 stack, stamp-cap, optional Helmholtz.

5.4. Audio

Standalone audio: `POST /analyze-audio`. File upload on every plan. Browse soundtrack (`live_tab`) is Premium and Pro.

Windows: mono 16 kHz, about 4 s each, center crop when needed, up to two windows, aggregate with **max**. Prewindowed clips may single-pass. Product video with usable audio can seat Helmholtz beside Muybridge (Covenant). That fused path has no public claim table until a measured fuse holdout clears. If the client sent audio and the windows fail: HTTP 503 — not a silent frames-only success.

5.5. Documents

PDF/DOCX inherit the text and image engines: extracted text through the text path; a capped set of page images through the image path; embedded clips through video. Purkyne Gate still runs on raw embeds. Clear errors: `text_too_short`, `page_limit`.

5.6. Confirm

A confirm is a retest the user chooses after a product score already exists. Sightengine (media) or Winston (text). First-party only (Check, Android, extension, Dashboard) — not the Pro API. 300 pooled units per UTC calendar month. Video and GIF cannot be confirmed. Confirm never writes a Purkyne stamp. Vendor results are commercial second opinions — not the in-house tables below.

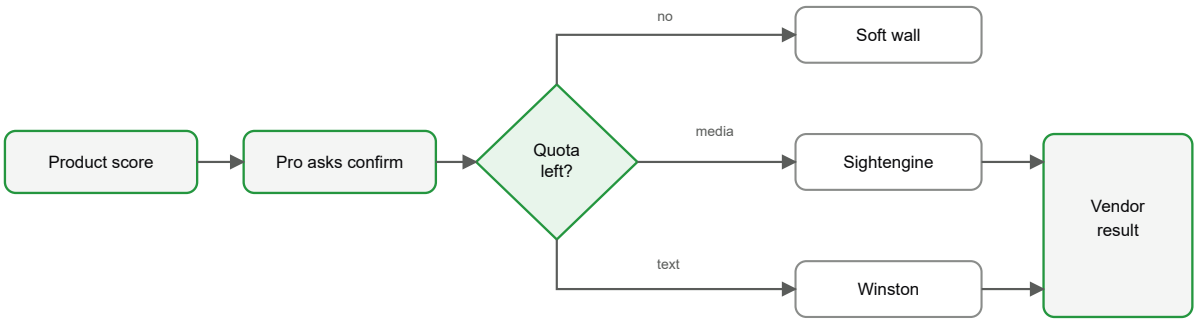

Figure 6 — Confirm: explicit retest after a product score; vendor results are separate from the gate tables.

6. Engine matrix

Public names are labels. Wire ids are API and cache keys. Purkyne Gate and Covenant are Express product checks, not rows here.

SLOT	ENGINE	ARTIFACT / STACK
Image	Vermeer@14.2 (`inhouse@14`)	Niepce 13.14 + Janssen 6.6 · `siglip2_base_patch16_224_linear_head_v14_2.npz`
Text EN	Valla@14.3 (inhouse-text@14)	Rask + Comenius + Sweet · `text_stack_head_v14_3_en.npz`
Text ZH	Valla (`inhouse-text-zh_v2`)	Ma Jianzhong · `text_stack_head_zh_v2_2.npz`
Text ES	Valla (`inhouse-text-es@14`)	Rask + Comenius + Bello · `text_stack_head_v14_3_es.npz`
Text PT	Valla (`inhouse-text-pt@14`)	Rask + Comenius + Ribeiro · `text_stack_head_v14_3_pt.npz`
Text FR	Valla (`inhouse-text-fr@14`)	Rask + Comenius + Littre · `text_stack_head_v14_3_fr.npz`
Text IT	Valla (`inhouse-text-it@14`)	Rask + Comenius + Ascoli · `text_stack_head_v14_3_it.npz`
Text DE	Valla (`inhouse-text-de@14`)	Rask + Comenius + Grimm · `text_stack_head_v14_3_de.npz`
Text NL	Valla (`inhouse-text-nl@14`)	Rask + Comenius + Vries · `text_stack_head_v14_3_nl.npz`
Video	Muybridge@4.4 (`inhouse-video@4`)	Plateau 3.8 + Blau 1.4 · `siglip2_base_patch16_224_linear_head_v4_4.npz` · stamp-cap
Audio	Helmholtz@3.1 (`inhouse-audio@3`)	Helmholtz · Dasheng-Base frozen · `dasheng_base_linear_head_audio_h3e.npz`

7. Image engine: Vermeer

Wire `inhouse@14`. Artifact `siglip2_base_patch16_224_linear_head_v14_2.npz`. Two SigLIP2 spines (Apache-2.0): 13.12 on transport, 6.6 on center. Linear heads Niepce 13.14 and Janssen 6.6 feed a two-logit stack (14.2). Serve on CPU via `/v1/analyze-pair` `kind=still`.

The stack npz is the combiner. Heads live beside it (`v13_14` Niepce, `v6_6` Janssen). Each view is embedded on its spine; the matching head yields a logit; the stack maps the pair to P(AI).

7.1. Holdout

Cut 0.5. Commercial-clean self-generated AI and diversified real photographs, with ordinary JPEG and resize stress.

CLASS	SET	N
Real	Wiki photographs	197
AI	Held-out generators	200

7.2. Gates

WHAT WE MEASURED	TARGET	RESULT	PASS
Average score on real photos	≤ 0.12	**0.025**	YES
Average score on AI images	≥ 0.88	**0.983**	YES
Gap between AI and real averages	≥ 0.55	**0.958**	YES
Balanced accuracy at our threshold	≥ 0.92	**0.9924**	YES

Cite the live 14.2 wire only.

8. Text engines: Valla

Eight language-specific heads. Latin seven share Rask and Comenius (loaded once) plus a named tiny and a small npz stack. Chinese is Ma Jianzhong (`inhouse-text-zh_v2`), independent of that trio.

8.1. English — Valla@14.3 (inhouse-text@14)

Artifact: `text_stack_head_v14_3_en.npz`. Architecture: Rask + Comenius + Sweet.
Human train: Gutenberg.org sentence windows (40%), Wikipedia (35%), WikiText (25%), sliced 520–1100 characters. AI train: Qwen2.5-7B, Mistral-7B, Grok-4.5. Holdout AI: Qwen2.5-1.5B (encyclopedic-heavy). Holdout human: equal-thirds of the same three sources. About 400 human and 400 AI samples.
Public English numbers are RAID-lite (Wikipedia blocks; news observed only). No published holdout from live ChatGPT or Claude scrapes.

8.2. English gates

WHAT WE MEASURED	TARGET	RESULT	PASS
Average score on human text	≤ 0.12	**0.029**	YES
Average score on AI text	≥ 0.85	**0.984**	YES
Balanced accuracy at our threshold	≥ 0.90	**0.979**	YES

Decision threshold 0.5. Cite RAID-lite.

8.3. Other languages

Same absolute floors as English (human mean ≤ 0.12, AI mean ≥ 0.85, balanced accuracy ≥ 0.90). Human samples: Wikipedia lead prose, continuous encyclopedic paragraphs. AI samples: open models, native-language encyclopedic prompts. Latin holdouts 200 human / 200 AI. Chinese holdout: Yi-1.5-6B + SmolLM2, encyclopedic, 400 human / 480 AI.

LANG	ARCHITECTURE	HUMAN MEAN	AI MEAN	BAL. ACC.
Chinese	Ma Jianzhong	0.047	0.922	0.954
Spanish	Rask + Comenius + Bello	0.029	0.993	0.982
French	Rask + Comenius + Littre	0.021	0.990	0.978
Portuguese	Rask + Comenius + Ribeiro	0.042	0.964	0.958
German	Rask + Comenius + Grimm	0.034	0.966	0.978
Italian	Rask + Comenius + Ascoli	0.016	0.943	0.972
Dutch	Rask + Comenius + Vries	0.056	0.963	0.950

Wires: `inhouse-text-zh_v2` · `inhouse-text-es@14` · `inhouse-text-fr@14` · `inhouse-text-pt@14` · `inhouse-text-de@14` · `inhouse-text-it@14` · `inhouse-text-nl@14`.

9. Video engine: Muybridge

Wire `inhouse-video@4`. Artifact `siglip2_base_patch16_224_linear_head_v4_4.npz`. Same two spines as stills; Plateau 3.8 (transport) and Blau 1.4 (center); 4.4 stack. Serve `/v1/analyze-pair` `kind=video`.
Still and video-frame distributions differ (motion blur, compression, generator fingerprints). Missing pair stack ⟶ clear error, not a silent still-head score.

9.1. Holdout

Fit: CogVideoX-2b and Wan2.1 (Apache-2.0), plus other commercial-clean families in the Muybridge mix. Holdout scoring uses the same pair2 stamps the product uses (n=2 inner frames on short clips). Product fold: max (stamp-cap). Cut 0.5.

CLASS	N	NOTES
Real	602	pair2 stamps; stamp-cap fold
AI	180	reserved generators in the Muybridge mix

9.2. Gates

WHAT WE MEASURED	TARGET	RESULT	PASS
Stamp-cap average on real bags	≤ 0.12	**0.022**	YES
Stamp-cap average on AI bags	≥ 0.85	**0.976**	YES
Balanced accuracy at our threshold	≥ 0.90	**0.9859**	YES

Cite the live 4.4 stamp-cap wire only.

10. Audio engine: Helmholtz

Wire `inhouse-audio@3`. Artifact `dasheng_base_linear_head_audio_h3e.npz`. Frozen encoder: Dasheng-Base (`mispeech/dasheng-base`, Apache-2.0). Head: logistic classifier we trained. Serve `/v1/analyze-audio`. Missing head ⟶ 503.

10.1. Holdout

Fit reals: LibriTTS + FSD50K (concat, no loop-tile). Fit AI: open speech and music generators (parler-mini, dasheng-audiogen) plus fal-family train and Imagine in the linear mix. Holdout speech: CodecFake and DFADD. Holdout music/SFX: AudioGen odd-seed partition. Holdout fal: prompt-id split. The public numbers are the overall holdout.

10.2. Gates

WHAT WE MEASURED	TARGET	RESULT	PASS
Average score on real audio	≤ 0.12	**0.018**	YES
Average score on AI audio	≥ 0.85	**0.989**	YES
Balanced accuracy at our threshold	≥ 0.90	**0.995**	YES

Cite the live Helmholtz holdout only. No fused-video accuracy table from Covenant seating.

11. Evaluation

For each media kind a test set is held out of training. We report:

- **Real / human mean** — should stay low.
- **AI mean** — should stay high on reserved generators.
- **Gap** — AI mean minus real/human mean, where the gate set includes it.
- **Balanced accuracy** — class-averaged accuracy at threshold 0.5.

An engine becomes the product default only after those checks pass. Holdout tables are **model** results. Purkyne Gate locks are a separate product check and are not mixed into these averages.

A hash cache keyed by `(clean_hash, engine)` may return a prior **product** score, including a stored Gate lock. That is operational, not a quality claim. The public DotCheck Index is the 30-day mean of first product scores (including Gate locks). It is not a substitute for the tables in this report.

12. Failure modes

CONDITION	BEHAVIOUR
Inference down or missing required head	HTTP 503. No silent Sightengine/Winston for Free or Premium.
Purkyne evidence missing or weak	Gate stays open. A named model scores. No invented 100% lock.
Pair-resized JPEG as Gate input	Never. Extract always runs on raw.
Unsupported language	Not scored.
Document too short / too many pages	`text_too_short` / `page_limit`.
Live-tab audio on Free	`403 live_audio_paid_only`. File audio remains.
Soundtrack sent and windows fail	HTTP 503. Not a silent frames-only success.
Share busy or disabled	`409 share_busy` / `503`. Share never fabricates a percent.
Confirm exhausted / not allowed	Soft wall. Video and GIF cannot be confirmed.

13. Limits

- Holdouts above only. New generators, heavy edits, screenshots-of-screenshots, and unusual compression can move **model** scores.
- Text: EN, ZH (Simplified), ES, PT, FR, IT, DE, NL. No published ChatGPT/Claude holdouts.
- Video is pair-frame stamps, not temporal forensics. Video holdout is reserved Wan clips, not a third unseen generator brand. GIF is video. No fused A/V claim table.
- Purkyne Gate can be fooled. Marks can be stripped, copied, or faked. Chat apps and feeds usually remove the print; then the models decide.
- Optional Pro Confirm (Sightengine/Winston) is not in these tables.
- Not proof of authorship. Not for employment, academic misconduct, credit, or similar automated decisions about a person.

14. Upstream models

Commercial-clean serve backbones and train/holdout generators. Hugging Face ancestor ids are not the product names.

MODEL	ROLE	LICENSE	USED BY
`google/siglip2-base-patch16-224`	image_video	Apache-2.0	Vermeer + Muybridge
`Oxidane/tmr-ai-text-detector`	text	MIT	Valla (Rask)
`intfloat/e5-small`	text	MIT	Valla (Comenius)
`prajjwal1/bert-tiny`	text	MIT	Valla (named tinies)
`fakespot-ai/roberta-base-ai-text-detection-v1`	text	Apache-2.0	Valla rollback @12
`mispeech/dasheng-base`	audio	Apache-2.0	Helmholtz
`hfl/chinese-macbert-base`	text	Apache-2.0	Valla (Ma Jianzhong)
`stabilityai/sd-turbo`	image fit	MIT	
`segmind/tiny-sd`	image fit	Apache-2.0	
`stabilityai/sdxl-turbo`	image fit	MIT	
`warp-ai/wuerstchen`	image fit	MIT	
`Qwen/Qwen2.5-7B-Instruct`	text fit	Apache-2.0	
`mistralai/Mistral-7B-Instruct-v0.3`	text fit	Apache-2.0	
`zai-org/CogVideoX-2b`	video fit	Apache-2.0	
`Wan-AI/Wan2.1-T2V-1.3B-Diffusers`	video fit	Apache-2.0	
`kandinsky-community/kandinsky-2-2-decoder`	image holdout	Apache-2.0	
`Qwen/Qwen2.5-1.5B-Instruct`	text holdout	Apache-2.0	
`Wan-AI/Wan2.1-T2V-1.3B-Diffusers`	video holdout	Apache-2.0	
`Wan-AI/Wan2.2-TI2V-5B-Diffusers`	video holdout	Apache-2.0	

Hub cards: https://huggingface.co/DotCheck · https://huggingface.co/collections/DotCheck/dotcheck-engines-6a659e57f7934cb510677785

15. Related reading

- **Model Card** — holdout numbers on the same docs page
- **How** — dotcheck.ai/how
- **Privacy Policy** — dotcheck.ai/privacy-policy
- https://dotcheck.ai/docs

Docs v2026.10 · 3 September 2026