CoreSR Inc.
c/o LaBarge Weinstein LLP, 800-515 Legget Drive, Kanata, ON K2K 3G4, Canada
help@coresr.ai · coresr.ai

PUB-007 · Version 1.2
Effective: 2026-11-17
Public Procurement & Governance

Model Version Pinning & Release Changelog

Public Endpoint: https://coresr.ai/api/public/model-changelog
Methodology & Reproducibility Contact: help@coresr.ai · security@coresr.ai

Systematic reviews, meta-analyses, and clinical practice guidelines must be **methodologically reproducible**: an author team submitting a PRISMA 2020 manuscript or updating a living guideline needs to report the exact model generation, training snapshot date, and calibration operating point that scored each citation or proposed each appraisal item.

CoreSR pins exact model release identifiers onto every persisted AI record and includes them in Screening Audit CSVs and One-Click Full Account ZIP archives (Control 4.4).

1. Active Production Model Matrix (as of November 15, 2026)

Model Family	Pinned Release Tag	Serving Scope	Exported Provenance Fields	Artifact Integrity Digest (SHA-256)
TITAN-SR v6	`v6-2026-10-15`	Paid subscriptions & active seats (GPU)	`abstract_ensemble_version`, `abstract_sens_target_used`, `abstract_applied_threshold`	`eb0ac9c7...23a25304`
TITAN-SR dual3	`dual3-2026-07-12`	Free & 7-day Trial CPU screening tier	`abstract_ensemble_version`, `abstract_sens_target_used`, `abstract_applied_threshold`	`7ef195ea...53d9ef92`
TITAN RCT Pre-Filter	`rct-prefilter-v1`	Optional RCT design pre-filter across TITAN batches	`rct_probability`, `rct_flag`	`c02809ba...03dc70c9`
GLASS v3	`glass-v3` (6 specialist routes)	Grounded full-text screening, PICO/outcome extraction, Risk-of-Bias, GRADE writing & synthesis (65,536-token context)	`model_used`, `advisor_model_runs_json`	Promoted 2026-10-17; self-hosted CoreSR weights

Model Family	Pinned Release Tag	Serving Scope	Exported Provenance Fields	Artifact Integrity Digest (SHA-256)
ROB BERT v6.1	`v6.1`	Two-stage Cochrane RoB 2 reader + 22-head signalling-question checker (8,192-token context)	`bert_model_version`	Promoted 2026-10-17; paired with deterministic Sterne 2019 rollups

> **Batch Generation Isolation Invariant:** CoreSR prevents mixing calibrated scores from two different TITAN-SR ensemble generations inside a single screening batch unless the reviewer explicitly initiates a full batch re-screen or reclassification.

---

2. Dated Production Release Log

2026-11-17 — Kaplan–Meier Figure Reading (KM Cor) Uses the External Vision Model on Every Plan

- **Models Affected:** KM Cor figure reading only (Google Gemini vision model). `TITAN-SR`, `GLASS v3` and `ROB BERT v6.1` are unchanged.
- **Summary of Changes:**
  - KM Cor now reads Kaplan–Meier figures with the external vision model in **Team** and **Institution** workspaces, as it already did on other plans. When a reviewer runs **KM Cor** on a Kaplan–Meier figure, the cropped image of that figure, its caption, and the review's subgroup names are sent to **Google's Gemini vision model** through a commercial model API, in Team and Institution workspaces as on every other plan. Nothing else is sent: not the PDF, not any other page, and no screening, extraction, risk-of-bias or GRADE data. KM Cor runs only when a reviewer starts it on a figure, and a figure can instead be traced by hand, which uses no AI.
  - Every other AI workflow in Team and Institution workspaces remains blocked in server code from calling external commercial LLM APIs.
- **Methodological Impact:** KM Cor was unavailable in Team and Institution workspaces; it now returns the same reconstruction as on other plans. The reader recorded on each digitized figure identifies the model used.

2026-11-14 — Zero Vendor Egress & Export Provenance Hardening

- **Models Affected:** All production AI workflows (`TITAN-SR`, `GLASS v3`, `ROB BERT v6.1`, CoreGRADE Advisor).
- **Summary of Changes:**
  - Enforced server-level zero third-party commercial LLM egress across all **Team** and **Institution** workspaces, ensuring Risk-of-Bias, extraction, screening, and synthesis pipelines never fall back to external commercial LLM APIs.
  - Expanded One-Click Full Account ZIP Exports (`decisions.csv` and `coregrade_assessments.csv`) to include pinned `abstract_ensemble_version`, `abstract_sens_target_used`, `abstract_applied_threshold`, and `advisor_model_runs_json` provenance fields alongside existing Screening Audit CSV exports.
- **Methodological Impact:** Zero change to model weights, scores, or calibration thresholds; strengthens export auditability and institutional data-boundary guarantees.

2026-10-17 — Promotion of GLASS v3 (Six Specialist Routes, 65,536-Token Context) & ROB BERT v6.1

- **Models Affected:** `GLASS v3` and `ROB BERT v6.1`.
- **Summary of Changes:**
  - Promoted **GLASS v3** across all six grounded synthesis routes (`glass-v3-screen`, `glass-v3-extract`, `glass-v3-rob2`, `glass-v3-rob-multi`, `glass-v3-grade-writer`, `glass-v3-agentic`) with a 65,536-token context

window, pre-adapted on 3,097,955 screening records and aligned for structured extraction, multi-instrument Risk-of-Bias, and GRADE certainty drafting.
    - Promoted **ROB BERT v6.1** (4,237 deduplicated training worksheets, 532 held-out test worksheets, 8,192-token checker context, calibrated confidence gating across 22 Cochrane RoB 2 signalling heads).
- **Methodological Impact:** Expanded single-pass full-text study reading to 65,536 tokens with strict verbatim source-quote verification; existing saved extractions, RoB assessments, and GRADE tables remain unchanged unless explicitly re-run by a reviewer.

2026-10-15 — Promotion of TITAN-SR `v6-2026-10-15` (3-Member Calibrated Rank-Mean Ensemble)

- **Model Affected:** `TITAN-SR v6-2026-10-15`.
- **Training & Calibration Snapshot (2026-10-14):**
    - Trained on 3,097,955 bibliographic records (315,052 final-inclusion positives) across 19,926 systematic-review groups, targeting **final inclusion in the published review**.
    - Three complementary biomedical transformer encoders combined via Platt scaling and empirical CDF rank-mean calibration.
    - Operating-point recall thresholds (`0.995` default, `0.99`, `0.97`, `0.96`, `0.95`, `0.93`) anchored at the conservative 5th percentile of review-cluster bootstrap replicates across 80 held-out calibration reviews and evaluated on 60 unseen test reviews (91,538 records; 77.98% validation specificity at 99.5% target recall, 95.50% at 95.0% target recall).
- **Methodological Impact:** Promoted as the primary GPU screening ensemble for paid subscriptions while retaining `dual3-2026-07-12` for Free/Trial CPU screening. Existing batches retain their recorded `ensemble_version`.

2026-07-12 — Release of TITAN-SR `dual3-2026-07-12`

- **Model Affected:** `TITAN-SR dual3-2026-07-12`.
- **Summary of Changes:** Introduced delivered-recall shift-margin calibration (`0.995` – `0.93` sensitivity targets) across a 3-member biomedical encoder ensemble. Remains in active production as the dedicated CPU screening engine for Free accounts and 7-day trials.

---

3. Recommended Manuscript Citation Format for Authors

Reviewers reporting AI-assisted workflows under **PRISMA 2020** or **PRISMA-trAIce** may cite the exact pinned release recorded in their exported Screening Audit CSV:

> *"Title and abstract screening was assisted by CoreSR TITAN-SR (`ensemble_version: v6-2026-10-15` , sensitivity target `0.995`), and structured full-text extraction and Risk-of-Bias proposals were generated using CoreSR GLASS (`v3` , verbatim-quote grounded), with qualified human reviewer verification of all final inclusion decisions, extracted values, and domain judgements."*

ISSUED ON BEHALF OF CORESR INC.

Tyler Pitre, MD PhD FRCPC
Founder and Chief Executive Officer, CoreSR Inc.
tyler.pitre@coresr.ai · coresr.ai