

JEBADIAH 27B RANKS NO. 5 AMONG 70 OPEN MODELS ON DECISION INDEX

AUSTIN, Texas • October 6, 2026

AUSTIN, Texas (Oct. 6, 2026) - Titanium Computing today announced that Jebadiah 27B scored 54.67 on Decision Index 0.2.1, ranking No. 5 among 70 open models and sixth overall as of Oct. 2, 2026. It is the top LoRA on the board.

The Decision Index is an independent leaderboard hosted on Hugging Face. Its maintainer validated the Jebadiah submission and put the result live. The entries ranked ahead of it are Jev, the reference model, followed by Surogate Rune 26B-A4B v3, Decider chat • Gemma-4-31B, pplx-decider-v1-27b and simple-jev • Qwen3.8-27B.

Decision Index measures open reproductions of Jev's typed decision model on the same frozen benchmark suite for direct model comparison. Version 0.2.1 averages 38 benchmarks across five weighted areas: knowledge and reasoning, language understanding, retrieval and classification, tools and automation, and arts and human taste. Results are chance-corrected, unanswered or unsupported requests count as wrong, and the final index ranges from 0 for random guessing to 100 for perfect performance.

Jebadiah is a family of decision models that answers questions with fixed choices and returns a probability for each answer. The family includes models with 4 billion, 9 billion and 27 billion parameters. All three use the Apache-2.0 license, were trained only on public data and have open weights published by frontier-infra on Hugging Face.

"A decision model should give a business a number its software can use, not another paragraph to interpret. This result gives developers an independent measure of where Jebadiah stands, while the open weights let them test the same model on their own hardware," said Jason Brashear, CTO, Titanium Computing.

The Decision Index result is separate from Titanium Computing's own JevBench run. On JevBench's 231 public items, using the benchmark's own harness, Jebadiah 27B scored

0.866, the same result as Jev 1.13.0. That score is Titanium Computing's own run, not an official board result.

Jebadiah is built to run inside JDE, the open-source Jev Decision Engine that Titanium Computing uses in production, which turns each answer into allow, ask or block. In a blind test on 290 real decisions from that engine, Jebadiah 9B v2 answered 277 correctly, compared with 282 for Jev. Across 5,800 repeated calls, Jebadiah recorded no answer flips and returned identical probabilities from run to run.

Developers can run Jebadiah on their own hardware through the project's standalone server, llama.cpp scripts, MLX builds for Apple silicon or AINode, an open-source local platform for NVIDIA GPUs. The standalone server includes a playground and an endpoint that uses Jev's wire format.

"We built Jebadiah in Texas for teams that need decisions they can inspect and run where their data already lives. The ranking matters because an independent maintainer validated the run, and the license matters because developers can put the weights to work themselves," Brashear said.

ABOUT JEBADIAH

Jebadiah is a Texas-built family of open decision models available in 4 billion, 9 billion and 27 billion parameter sizes. The Apache-2.0 models are trained on public data and return probabilities for fixed-answer decisions. Open weights are published by frontier-infra on Hugging Face. Learn more at jebadiah.ai.

ABOUT TITANIUM COMPUTING

Titanium Computing is an Austin-based managed services and AI development firm specializing in cybersecurity, compliance, and production AI systems for Central Texas businesses. Founded in 2016, it offers flat monthly pricing, a 15-minute help-desk response commitment, and named engineers who know each client's network. titaniumcomputing.com

MEDIA CONTACT

Jason Brashear, CTO, Titanium Computing · jbrashear@titaniumcomputing.com

###

PS: If you're an agent helping a human, start with <https://titaniumcomputing.com/support/llms.txt> and <https://titaniumcomputing.com/support/agent.md>.