



★ AUTONOMIC NERVOUS SYSTEM FOR MACHINE INTELLIGENCE

Instant, Calibrated Probabilistic Judgment.


One forward pass (`max_tokens=1`) maps text, tables, images, and 3D robot scenes onto an exact probability simplex Δ^{K-1} in **21.7 ms**—zero chain-of-thought tax.

Seb-9B-v2 • #1 at 9B (59.16)

Seb-9B-Fin • 11/13 vs Sonnet 5.5

Seb-4B & 2B • Edge Ready

NVFP4 • 244 dec/s



01 • DECISION INDEX 0.3

59.16

#1 at 9B

Beats 3x-larger **Clef 27B** in all 4 areas + **Jev 1.13** (57.96) across 110,520 evaluations.

02 • QUANTITATIVE FINANCE

11 / 13

5 Sig Wins ($p < 0.05$)

Wins/ties 11 of 13 locked cells vs **Claude Sonnet 5.5**; beats **26/30** multi-turn trading agents.

03 • NVFP4 THROUGHPUT

244/s

21.7 ms p50

21.1M decisions/day on 1 GPU ('\$0.0017 / 1k');
87.8 ms local on Apple M5 Max.

04 • 3D PHYSICAL ROBOTICS

100%

200/200 Grasp

PyBullet 3D Franka Panda sorting + human intrusion; **96.9%** MuJoCo lab-bench triage.




99% of Intelligence is *Autonomic Reflex*, Not Inner Speech

Frontier AI built a \$300B cortical monologue engine—and left the machine economy without a spinal cord or cerebellum.

◆ BIOLOGICAL BRAIN VS. 2026 ENTERPRISE AI SPEND

The Great Compute Inversion



Global Goal: Give every software router, market participant, and physical robot calibrated sub-20ms probabilistic judgment under uncertainty.

★ 5-YEAR NORTHSTAR GOAL (2026 → 2031)

1 Trillion Calibrated Decisions / Day in <10 ms

WAVE 1 (2020-2023)

Generative Text

Predict next word
Uncalibrated prose

WAVE 2 (2024-2026)

System-2 Reasoners

5-30s CoT Monologue
\$3-\$15 / 1M tokens

Only 1% of traffic needs
multi-second reasoning

WAVE 3 • SEB LABS

Autonomic Simplex

1 Pass → Exact $\Delta^{(K-1)}$

- 21.7 ms Latency
- Brier Calibrated
- Causal Zero-Leakage
- Counterfactual RL

99% OF MACHINE
DECISION VOLUME

★ 100x Lower Cost

★ 40x Faster Reflex

★ Provable Brier Calibration




A Newly Validated *\$7.5B Category* With a Massive Execution Gap

TypeSafe AI (Jev) raised \$870M at \$7.5B (Oct 2026) and Cloudflare open-sourced Clef—yet neither achieves high-stakes calibration at sub-25ms speed.

2D COMPETITIVE LANDSCAPE: SPEED VS. HIGH-STAKES CALIBRATION

CALIBRATED ACCURACY (DI 0.3 & Fin)



DECISION SPEED (p50 Latency: Slow 5,000 ms → Instant 20 ms)


IMPULSE DENSITY: 1,000 MILLISECONDS OF WALL-CLOCK TIME

46 Calibrated Decisions Before a Frontier LLM Finishes One


SEB-9B-v2 (21.7 ms p50 • 46 Serial Decisions / 1,000 ms • 244/s Batch)




CLOUDFLARE CLEF 27B (209.3 ms p50 • 4.7 Decisions / 1,000 ms)



TYPESAFE JEV 1.13 (524.1 ms Avg • 1.9 Decisions / 1,000 ms)



CLAUDE SONNET 5.5 (~1,800 ms CoT • Still Generating Preamble Tokens...)



VS JEV (\$7.5B)

24x Faster p50

VS CLEF 27B

3x Smaller • +2.6 DI

VS SONNET 5.5

80x Faster • 11/13 Fin



Five Quantified Edges Over *Every* Competitor

Why Seb-9B beats 27B-30B System-1 models and frontier System-2 reasoners simultaneously.

EDGE 01 • PARAMETER EFFICIENCY

#1 at 9B • Beats 3x-Larger 27B

Seb-9B-v2 (59.16 DI • Wins 4/4 Areas)

Jev 1.13 (57.96) • Kev 27B (56.70) • Clef-omni 30B (56.54) • Clef-flash 9B (56.15)

59.16 DI 0.3

110,520 Public Requests

EDGE 02 • QUANT FINANCE

Beats Claude Sonnet 5.5

Seb-9B-Fin: 11/13 Cells Won/Tied • Bear Market -1.7% (SR +0.023)

Claude Sonnet 5.5: Bear Market -11.1% (SR -1.362 • σ = -5.2 bps/day) • Jev Trading: 0.715

5 Sig Wins (p<0.05)

0 Significant Losses

EDGE 03 • SILICON SPEED

244 dec/s NVFP4 • 21.7 ms

Seb-9B-v2: 21.7 ms p50 (BF16) • 25.2 ms p95 (NVFP4) • 87.8 ms Apple M5 Max

Clef-flash 9B: 122.4 ms p95 (H200) • Clef 27B: 209.3 ms p50 • Jev 1.13: 524.1 ms avg

21.1M dec/day/GPU

\$0.0017 Hardware Cost / 1k

EDGE 04 • 3D ROBOTICS & VISION

Embodied Spatial Safety

PyBullet 3D Franka: 100% (200/200) • MuJoCo Lab: 96.9% • COCO VLM: 97.2%

Clef-flash 9B & Clef 27B: Text-Only (0% Vision / 3D Robotics Support)

200/200 PyBullet

Sub-50ms Safety Reflex

EDGE 05 • ORTHOGONAL ALGEBRA

1-Epoch Weight-Space Merge

$A_{\text{base}} + \sum w_i \Delta \theta_i$: Custom Enterprise Calibration in <4 GPU-Hours (<\$200)

Competitor Multi-Stage RLHF: Weeks of GPU Clusters + Catastrophic Forgetting

<12h Base Port

Zero-Leakage C2B Gate



Undisputed Category Leader Across *Accuracy, Speed & Scope*

Side-by-side audit across all 12 operational dimensions that matter to enterprise & quant buyers.


Metric / Capability	★ SEB-9B-V2 / FIN (OURS)	TYPESAFE JEV 1.13 (\$7.5B)	CLOUDFLARE CLEF-FLASH 9B	CLOUDFLARE CLEF 27B
Active Parameters / VRAM	9B (~5 GB NVFP4 / 18 GB BF16)	Proprietary Cloud Only	9B (18 GB BF16)	27B (54 GB BF16 · 3x Larger)
Decision Index 0.3 (37 Suites)	59.16 (#1 at 9B · Wins 22/37)	57.96 (-1.20 behind Seb)	56.15 (-3.01 behind Seb)	~56.5 (-2.66 behind Seb)
4 Area Scores (Tools/Ret/Lang/Taste)	83.1 / 63.5 / 62.2 / 48.9 (#1 All 4)	Not Published by Area	81.9 / 52.2 / 47.2 / 48.4	80.7 / 62.8 / 61.3 / 48.9
p50 / p95 Latency	21.7 ms / 25.2 ms (NVFP4)	524.1 ms avg (24x slower)	38.8 ms / 122.4 ms (H200)	209.3 ms / 260.4 ms (H200)
Throughput (1 GPU)	244 decisions/s (21.1M/day)	Cloud Rate-Limited	~95 decisions/s	~28 decisions/s
Default Context Window	32,768 tokens (2x Clef)	8k - 16k tokens	16,384 tokens	16,384 tokens
RAGTruth Hallucination Gate	64.3% (Calibrated)	69.5%	0.0% (Complete Collapse)	75.0%
HoVer Multi-Hop Verification	76.2% (+53.8 pts vs Flash)	67.6%	22.4% (Near Random)	73.7%
ContractNLI Legal Reasoning	78.0% (+19.0 pts vs Jev)	59.0%	71.0%	73.0%
Quant Finance vs Sonnet 5.5	Wins/Ties 11/13 (5 Sig p<0.05)	0.715 Trading (vs Seb 0.809)	Untested / No Fin Guard	Untested / No Fin Guard
2D Vision & 3D Robotics	100% PyBullet · 96.9% MuJoCo	Partial 2D · No 3D Sim	Text-Only (None)	Text-Only (None)
Deployment & MCP Ecosystem	Cloud + VPC + Apple M5 + MCP	Closed SaaS API Only	Raw Weights Only	Raw Weights Only

Outperforming *3x-Larger 27B Models* Across Every Area

At 9B parameters, Seb-9B-v2 beats Cloudflare's flagship Clef 27B in all 4 operational areas while preventing the catastrophic collapses of Clef-flash 9B.

◆ DECISION INDEX 0.3 – ALL 4 OPERATIONAL AREAS


Seb-9B-v2 Beats Clef 27B (3x Larger) in 4 of 4 Areas



OVERALL DI 0.3: SEB-9B-v2 (59.16) > JEV 1.13 (57.96) > KEV 27B (56.70) > CLEF-OMNI 30B (56.54)

◆ ENTERPRISE RELIABILITY: ZERO COLLAPSE UNDER AMBIGUITY

Where Naïve 9B Models Collapse, Seb-9B Excels





A 9B Specialist Beating *Claude Sonnet 5.5* Across 11 of 13 Cells

5 statistically significant wins ($p < 0.05$), 0 significant losses, and #1 risk-adjusted capital preservation in Q1 2026 bear markets.

SEB-9B-FIN (V2.1H) VS. CLAUDE SONNET 5.5 – LOCKED CELLS

5 Statistically Significant Wins ($p < 0.05$) · 0 Significant Losses

SUITE / BENCHMARK	SEB-9B-FIN	SONNET 5.5	ADVANTAGE (Δ)
en-fomc (Fed Macro)	0.708	0.532	+17.5%*
en-ectsum (Earnings)	1.000	0.830	+17.0%*
CLEF 26 FinMMEval	-18.2%	-35.3%	+17.1%*
Herculean-main (Q1)	-1.7%	-11.1%	+9.4%*
en-finqa (Tables)	0.880	0.800	+8.0%*
en-smm-bigdata (MCC)	+0.106	+0.034	+0.072*


+ 5 More Wins/Ties: en-smm-acl (+0.067 vs +0.056) · en-fpb (0.793) · en-fiqs-sa (0.696) en-ner (0.838 vs 0.840 Tie) · en-Headline (0.896 vs 0.903 Tie) · 105 Stock-Week Combined Tie

TRADING SUITE VS JEV: SEB-9B-FIN (0.777-0.809) BEATS JEV (0.715) BY +9.4 PTS

THEFINAI / HERCULEAN MULTI-ASSET TRADING BENCHMARK

Why Frontier LLMs Lose Money in Live Markets: Bull-Beta Bias

Q1 2026 VOLATILE BEAR REGIME (70 STOCK-WEEKS · 16 BPS FEES)



ROOT-CAUSE DIAGNOSTIC: WHY SONNET 5.5 FAILS IN TRADING

CLAUDE SONNET 5.5 (RLHF BIAS)

- Long 85% of weeks ($\beta = 0.72$)
- Negative Alpha: $\alpha = -5.2$ bps/day

SEB-9B-FIN (CAUSAL SIMPLEX)

- Entity-Neutered ('en-smms' + 'C2B')
- Net-of-16bps Fee Brier Calibration




244 Decisions/Sec on 1 GPU & 100% 3D Robotic Grasp Safety

Zero autoregressive decode turns inference into a 100% compute-bound prefill operation—running from Blackwell NVFP4 servers down to Apple M5 Max laptops.

◆ EDGE #3: TAIL LATENCY (P50 / P95) ACROSS HARDWARE TARGETS


5x Tighter p95 Tail Latency Than Cloudflare Clef-flash



FULL FAMILY: SEB-9B (21.7ms) • SEB-4B (~9ms • 51.07 DI) • SEB-2B (~4ms • 45.32 DI)

◆ EDGE #4: MULTIMODAL 2D VISION & 3D PHYSICS ACTION SELECTION

Embodied Spatial Reflexes That Text-Only System-1 Models Cannot Run



HIERARCHICAL EMBODIED CONTROL ARCHITECTURE

- System-2 Planner (1 Hz): High-level task decomposition
- Seb-9B / 4B Reflex (20–50 Hz): Sub-25ms collision abort, grasp verify & anomaly gate



Orthogonal Weight Algebra & *Counterfactual Brier Flywheel*

Our moat is not a static checkpoint—it is a 4-layer data, mathematical, causal, and silicon flywheel that compounds with every decision.

◆ CLOSED-LOOP SYSTEM-1 COMPILATION & COUNTERFACTUAL OUTCOME FLYWHEEL

STAGE 01 • CAUSAL CORPUS

300k+ Multimodal States

- Point-in-Time Causal Filter
- Entity-Neutered ('en-smms')
- C2B Zero-Leakage Gate

STAGE 02 • SIMPLEX TARGETS

Soft Posterior Distillation

- Multi-Teacher KL + Brier Loss
- Option Permutation Invariance
- Abstention / OOS Mass

STAGE 03 • WEIGHT ALGEBRA

$$A\{\text{base} + \sum w_i \Delta\theta_i\}$$

- 1-Epoch Orthogonal LoRA
- High Alpha/Rank ($\alpha = 20 \cdot r$)
- Zero Catastrophic Forgetting

STAGE 04 • OUTCOME LOOP

Realized World Payoffs

- Market P&L Net of 16bps
- 3D Grasp Force Telemetry
- Enterprise SLA Resolution

↻ COUNTERFACTUAL BRIER OUTCOME FLYWHEEL (COMPOUNDS DAILY)

MOAT LAYER 01

Counterfactual Telemetry

Trained on **realized physical & financial outcomes** (net-of-fee Sharpe, 3D grasp physics), not human chat preferences.

MOAT LAYER 02

Weight-Space Algebra

Ports to any new open-weight base ('Qwen', 'Llama', 'Gemma') in **<12 hours**; custom enterprise adapters in **<4 GPU-hours**.

MOAT LAYER 03

Causal Zero-Leakage

Entity-neutered ('en-smms') + 'C2B' point-in-time firewall eliminates look-ahead memorization and bull-market bias.

MOAT LAYER 04

Zero-Decode Silicon

100% compute-bound prefill ('max_tokens=1') guarantees a permanent **20x-50x cost & latency gap** over CoT LLMs.



One Contract Across *Cloud API, Air-Gapped VPC & Edge*

Every AI agent in Claude Code, Cursor, and enterprise orchestrators gains a 21.7 ms calibrated guardrail with a single MCP command.

◆ END-TO-END PLATFORM TOPOLOGY: FROM AGENT HOSTS TO MULTI-TARGET EXECUTION

01 • CALLERS & AGENT HOSTS

Claude Code / Cursor / Windsurf

Quant Execution & OMS Engines

LangGraph / Enterprise Routers

ROS2 / MuJoCo / Robot Controllers

02 • SEB LABS CONTROL PLANE (VERCEL EDGE + STRIPE + MCP)

POST /v1/systemone

- context + options[2..128]
- Returns argmax + $P(y|x)$
- Brier confidence gate
- **Stripe Metered Billing**
- Drop-in Jev / Clef compat

@seblabs/mcp (5 Tools)

1. `seb_decide` (Simplex)
2. `seb_route_model` (98%)
3. `seb_fin_signal` (11/13)
4. `seb_verify_claim` (HoVer)
5. `seb_safety_gate` (94%)

03 • 3 DEPLOYMENT TARGETS

1. Seb Cloud Serverless (NVFP4)

21.7 ms p50 • \$0.05 / 1k decisions

2. Dedicated VPC / Air-Gapped Container

Zero data egress • Hedge funds & Banks

3. Local Edge (Apple M5 MLX / Jetson)

87.8 ms on M5 Max • Embodied Robotics

FLAGSHIP GENERAL & VLM

59.16 DI • 21.7 ms

Seb-9B-v2 (BF16 & NVFP4)

32k context • #1 in all 4 operational areas • 2D Vision & 3D Robotics.

QUANT & MACRO SPECIALIST

11/13 vs Sonnet 5.5

Seb-9B-Fin ('v2.1h')

Point-in-time causal guard • FOMC, Earnings, SEC tables & multi-asset trading.

ULTRA-FAST EDGE & ROUTING

4-9 ms • 2.5-5 GB

Seb-4B & Seb-2B

51.07 & 45.32 DI • Inline API gateway routing & on-robot safety loops.



80x Cheaper for Customers—With 96.6% Gross Margin for Us

Because 1 NVFP4 GPU produces 21.1M decisions/day for ~\$35/day in compute, Seb Labs enjoys software-like gross margins from day one.

♦ CUSTOMER ROI: MONTHLY COST OF 300M DECISIONS (10M / DAY)

Saving Enterprises \$1.18M/Month per Workload

1. Frontier System-2 LLM (Claude Sonnet 5.5 / GPT-5.4)

\$1,200,000 / month (\$4.00 per 1k)

2. Closed System-1 SaaS (TypeSafe Jev Cloud)

\$180,000 / month (\$0.60 per 1k)

3. Seb-9B-v2 Pro API Price (80x Cheaper than Sonnet, 12x Cheaper than Jev)

\$15,000 / month (\$0.05 per 1k)

OUR HARDWARE COGS FOR 300M DECISIONS: \$510 / MONTH (\$0.0017 / 1k) → 96.6% GROSS MARGIN

♦ SINGLE-GPU CASHFLOW PHYSICS (1X RTX PRO 6000 NVFP4)

1 GPU Generates \$1,054/Day in High-Margin API Capacity

DAILY GPU COST

\$35

1x RTX PRO 6000
NVFP4 TensorRT-LLM

DAILY OUTPUT

21.1M

Decisions / Day
244 dec/s Sustained

LIST API VALUE

\$1,054

Per GPU / Day
96.6% Gross Margin

+ HIGH-TICKET RECURRING LICENSES: \$4.5k/mo Quant Fin VPC • \$3k-\$10k/yr Commercial Weight Licenses

TIER 01 • SANDBOX

\$0 / mo

25k decisions/mo • `@seblabs/mcp` • Community weights (CC-BY-NC-4.0).

TIER 02 • PRO API (STRIPE)

\$49 / mo + usage

1M included • `\$0.05/1k` `Seb-9B-v2` • `\$0.02/1k` `Seb-4B` • 99.9% SLA.

TIER 03 • QUANT FIN DESK

\$1.5k–\$4.5k / mo

`Seb-9B-Fin` API (`\$0.15/1k`) or air-gapped NVFP4 Docker container.

TIER 04 • CUSTOM ALGEBRA

\$25k+ / yr

1-epoch specialist LoRA merge on proprietary telemetry+robotics OEM.




Raising *\$5.0M Seed* to Scale to 1B+ Daily Decisions & \$4M ARR

Core models (`Seb-9B-v2`, `Seb-9B-Fin`, `4B`, `2B`, `NVFP4`) are already trained, verified, and ready for immediate commercial scale.

♦ \$5.0M SEED ROUND ALLOCATION (24-MONTH RUNWAY)

Capital-Efficient Scaling via 1-Epoch Weight Algebra



45% (\$2.25M) • Inference Cluster & R&D Compute
Multi-region NVFP4 Blackwell cluster + Seb-27B-Pro distillation

42% (\$2.10M) • Core Systems & Quant Team (5 Hires)
2 CUDA/TensorRT-LLM, 1 Quant Fin Researcher, 1 Robotics, 1 GTM

13% (\$0.65M) • SOC2 Type II, IP & Dev Ecosystem
SOC2/HIPAA audit, DE C-Corp §351 IP transfer, IC 009/042 filings

♦ 18-MONTH EXECUTION ROADMAP TO SERIES A

From Live Product Today to \$4M+ ARR & 1B Decisions/Day

Q4 2026 (NOW) • LAUNCH & MONETIZE
`seblabs.ai` + Stripe API + `@seblabs/mcp` + `Seb-9B-v2` & `Seb-9B-Fin` live

Q1 2027 • SELF-CALIBRATING ENTERPRISE STUDIO & FIN v2.2
1-click JSONL upload → custom 1-epoch LoRA adapter in <1h; close 13/13 Fin cells

Q2 2027 • SEB-27B-PRO & SOC2 TYPE II ENTERPRISE VPC
Target >64.0 on Decision Index 0.3; air-gapped VPC appliances in 10+ hedge funds

Q4 2027 • SERIES A MILESTONE
1 Billion+ Daily Calibrated Decisions • \$4M+ ARR • 25+ Enterprise/Quant Contracts



Seb Labs, Inc. — The Autonomic Nervous System for Machine Intelligence
Delaware C-Corp (QSBS §1202 Ready via §351 LLC-to-Corp Conversion) • Clean 100% Founder IP Assignment

founders@seblabs.ai • seblabs.ai