AEVION QVenture — Investment Memo

Generated 2026-07-11 · AEVION AI Investment Analyst · not investment advice

ShelfSense

AI Applications (vertical SaaS) · seed · US · raising $3,500,000

Score 71/100 — WATCH (conviction: medium)

Investment memo

Verdict: watch, but participate small — this is a conviction-scaled toe-hold, not a lead check. ShelfSense has genuine early signal ($26k MRR, 12 CPG brands, reps 3x faster), and the single strongest reason to lean in is workflow lock-in: switching costs and proprietary planogram data can, if defended, elevate this above a commodity scanner in a 37%-CAGR market. The single strongest reason against is commoditization — foundation-model vision (GPT-4V/Gemini) is closing the CV gap while scaled incumbents (Trax, Standard AI) own the enterprise relationships, so value may never accrue above the model layer. On entry: write a $1,750,000 ticket (below the $2.04M target, holding to the exposure cap) for roughly 10% at the ~$16.9M pre, staged against two diligence gates — audited logo retention/NRR and contractual proof that scan data and model derivatives are not assigned to CPG customers. Reserve ~$3.06M for pro-rata. Pass if either gate fails; the concentration and data-ownership tails are dealbreakers, not haircuts.

Narrative engine: live model (anthropic)

Entry strategy

Ticket: $2,039,000 target (range $1,019,500–$1,750,000)
Target ownership: 10%
Valuation band (pre-money): $8,504,000 / $16,890,000 / $33,780,000
Return: 7.31x expected (16.9x base) · ~32.9% IRR over 7yr · loss prob 57%

Deployment schedule:

- 40% — Entry: On close, after founder + IP + cap-table diligence.
- 35% — Milestone: Product-market fit signal (retention cohort / first repeatable revenue).
- 25% — Pro-rata: Reserve for next priced round to defend ownership.

Portfolio: Size at ~2.2% of a diversified venture portfolio (fractional-Kelly, conviction-scaled). Reserve 3,058,500 USD for pro-rata follow-on.

Score breakdown

Market size & growth — 55/100 (weight 20%)
~$45B TAM, 37% CAGR (AI Applications (vertical SaaS)).

Timing / tailwinds — 100/100 (weight 10%)
Sector growth 37% vs. 12% neutral baseline.

Moat / defensibility — 74/100 (weight 15%)
Dominant defensibility here: switching costs.

Unit economics potential — 69/100 (weight 15%)
~70% mature gross margin, capital intensity 35%.

Team / execution signal — 68/100 (weight 12%)
revenue/customers cited

Scientific / tech feasibility — 100/100 (weight 10%)
agentic workflows, domain eval harnesses, retrieval + tool orchestration

Regulatory / legal headroom — 74/100 (weight 9%)
Regulatory intensity 40% (higher = more legal drag).

Competitive headroom — 41/100 (weight 9%)
Competitive intensity 85%. thin wrapper risk — value must accrue above the model layer.

Analyst council

Research Scientist — Retail shelf CV is technically mature; ShelfSense's edge is data/workflow, not model novelty — commodity risk is real

+ Core CV task (SKU detection, planogram matching from phone photos) is well-established: fine-grained object detection and few-shot recognition are solved to ~90-95% mAP on curated benchmarks; Trax, Standard AI, and Google's on-shelf availability APIs demonstrate feasibility. This is not a research frontier problem — the 100/100 feasibility score is justified but tells us little about defensibility

+ Real technical risk sits in the long tail: retail shelves have thousands of near-identical SKUs, packaging refreshes, occlusion, glare, and angle variance. Achieving >95% precision on out-of-stock/planogram calls in-the-wild (vs. curated demos) requires a large proprietary labeled image corpus — 60k shelves/month is a modest but compounding data moat if labeled and versioned well.

+ Defensibility must accrue above the model: domain eval harnesses (per-brand planogram schemas), reference-image management, and workflow integration into CPG field-force systems (Salesforce/Repsly/retail execution suites). The clipboard-replacement 3x speed claim is credible and the ROI is quantifiable, but the model layer is increasingly commoditized by foundation-model vision APIs.

+ Structural thin-wrapper risk is the central concern: incumbents (Trax raised >$1B, Standard AI, ParallelDots) and CPG-internal tools already do this at scale. Seed-stage $26k MRR / 12 brands is early validation but not yet evidence of a durable wedge against better-capitalized players.

! Accuracy ceiling in production: false OOS/compliance alerts erode rep trust fast; if in-wild precision lags demo numbers, the value prop collapses. Requires continuous, costly relabeling as SKUs/packaging change (retail SKU churn ~20-30%/yr).

! Commoditization: foundation-model vision (GPT-4V-class, Google Cloud Vision retail) is closing the capability gap, letting CPGs or SIs build in-house — value must live in proprietary planogram data and field-force workflow lock-in, which is not yet proven.

! Incumbent competition (85% intensity): Trax/Standard AI have scale, brand relationships, and data advantages; ShelfSense needs a defensible vertical wedge (e.g., underserved mid-market CPG or specific channel) to avoid being out-executed or acquired-for-talent.

Data Analyst — ShelfSense: real CPG traction at $26k MRR, but crowded retail-execution CV market caps defensibility

+ ARR ~$312k across 12 brands implies ~$26k ACV/brand — healthy for seed, but the 71/100 composite is dragged by 41/100 competitive headroom: Trax, Neurolabs, Wiser, and CPG in-house tools already sell shelf CV, so ShelfSense competes on execution not novelty.

+ TAM framing is inflated: $45B is all vertical AI SaaS, not retail-execution. Real SAM is US CPG field-audit software — realistically single-digit $B; SOM at seed is tens of thousands of enterprise field reps. Bottom-up TAM is MISSING.

+ Unit economics unverified: 70% mature gross margin assumes CV inference costs scale sub-linearly, but per-scan compute at 60k shelves/mo is undisclosed. CAC, payback, LTV, logo/net retention, and gross margin at current volume are ALL MISSING — the thesis lives or dies here.

+ Moat rated 74/100 on switching costs is plausible only if ShelfSense owns the planogram/compliance workflow and historical data; otherwise it's a thin wrapper on off-the-shelf detection models and value leaks to the model layer.

! Thin-wrapper/commoditization: retail-shelf detection is a near-solved CV problem; incumbents and CPGs can replicate core scan-to-alert, forcing price competition and eroding the 70% GM assumption.

! Enterprise sales concentration: 12 brands means high logo-concentration risk — losing 2-3 anchor accounts could halve MRR; no churn/retention data provided to assess durability.

! Land-expand ceiling: if ACV stays ~$26k and CPG procurement cycles are 6-12mo, $3.5M may not buy enough GTM runway to reach the ~$1M ARR / clean-cohort milestone a Series A requires.

Economist — ShelfSense: real wedge in retail execution, but CV commoditization caps durable rents at seed

+ Demand is inelastic where it matters: retail-execution/planogram compliance drives measurable sell-through; existing incumbents (Trax ~$1B+ raised, ImageAI/Repsly, GoSpotCheck) prove CPGs pay for shelf analytics, validating the $45B vertical-SaaS TAM though realistic SAM (US CPG field-audit software) is closer to $1-2B.

+ Moat is workflow lock-in and proprietary shelf-image data, NOT the CV model — 60k shelves/month builds a labeled dataset that compounds accuracy per SKU/retailer, the only credible defense against the 41/100 competitive-headroom score and thin-wrapper risk.

+ Unit economics are attractive (~70% mature GM) but current $26k MRR = ~$2,167/brand/month is subscale; land-and-expand into rep seats and per-shelf pricing must push ACV toward $50-100k to justify a vertical SaaS multiple.

+ Counter-argument (strongest bear case): Trax/incumbents already own tier-1 CPG relationships and phone-photo CV is now near-parity via foundation models — ShelfSense may be selling a cheaper commodity into a market where the buyer (large CPG) prefers a single global vendor, compressing pricing power.

! Commoditization: off-the-shelf multimodal models (GPT-4V/Gemini) collapse the CV moat; without proprietary planogram-matching + retailer-specific tuning, gross margin and pricing erode toward services.

! Customer concentration & procurement cycle: 12 brands, likely <$5k each — a single CPG churn is material, and enterprise CPG sales cycles (6-12mo) plus incumbent displacement risk threaten the growth needed for a Series A.

! Distribution disadvantage vs. Trax/incumbents who bundle shelf analytics into broader retail-execution suites, forcing ShelfSense into a feature rather than a platform and capping NRR.

Corporate & Regulatory Lawyer — ShelfSense: low-regulatory-drag CV SaaS; core legal risks are IP ownership, in-store image/privacy edge cases, and data-rights leverage from CPG customers.

+ Regulatory posture is favorable: B2B retail CV tool sits outside heavily-regulated verticals (no HIPAA/financial licensing). US privacy exposure is real but manageable — shelf photos incidentally capture shoppers/employees, triggering biometric statutes (Illinois BIPA $1k-$5k per violation, TX/WA analogs) and CCPA/CPRA if any facial data processed; recommend contractual ban on face capture + auto-

+ IP posture is the value question given 'thin wrapper' risk: if inference runs on third-party foundation/vision models, defensibility must come from proprietary planogram-labeled training data, eval harnesses, and CPG-specific fine-tuning. Diligence must confirm the company owns model weights/data pipeline and hasn't granted CPG customers ownership of derived data.

+ Data-rights are the strategic asset: 60k shelves/month of labeled retail imagery is the moat. Investor should require MSAs granting ShelfSense a perpetual, sublicensable license to use aggregated/anonymized scan data for model improvement — verify current 12 contracts don't cede this to brands.

+ Seed deal structure: standard priced round or capped SAFE with 1x non-participating liquidation preference, broad-based weighted-average anti-dilution, pro-rata rights, IP-assignment reps (all founder/contractor code assigned), and a specific rep on third-party model license compliance (OpenAI/Google TOS resale terms).

! IP fragility / wrapper dependency: if core CV runs on a licensed foundation model, a provider price hike, TOS change, or native competing feature (e.g., a retail-media giant shipping planogram detection) collapses the moat — the 41/100 competitive headroom and 85% intensity are the honest counter-argument that value may never accrue above the model layer.

! BIPA/biometric litigation tail: even incidental capture of identifiable faces in shelf photos exposes per-scan statutory damages in class-action-friendly states; at 60k scans/month this compounds fast if consent/blur controls aren't contractually and technically enforced.

! Customer-owned-data leakage: standard CPG procurement contracts often claim ownership of all data and outputs — if existing MSAs already assign scan data or trained model derivatives to brands, the company's central defensible asset is legally impaired at close.

Market data sources

- Grand View Research (2025) — Generative AI $22.2B in 2025, 37.6% CAGR to 2030
  https://www.grandviewresearch.com/industry-analysis/generative-ai-market-report
- ABI Research (2025) — Gen-AI software $63B in 2025 !' ~$220B by 2030 at ~29% CAGR
  https://www.abiresearch.com/blog/generative-ai-software-market-report-summary

Assumptions & limitations

- Market size / growth for AI Applications (vertical SaaS) is anchored to Grand View Research (2025): Generative AI $22.2B in 2025, 37.6% CAGR to 2030. Full citations are listed under "Market data sources".
- Stage norms reflect US-market seed deals; adjust for geography "US".
- Score is a screening signal, not a substitute for legal, financial, and technical due diligence.