AEVION QVenture — Investment Memo

Generated 2026-11-26 · AEVION AI Investment Analyst · not investment advice

FICTIONAL EXAMPLE: the company in this example is made up; any resemblance of names is coincidental. All figures are model output on invented inputs.

ShelfSense

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

Score 66/100 — WATCH (conviction: medium)

Signal coverage: 44% company-specific (1 metrics parsed from the plan); remainder from sector priors.

Investment memo

Verdict: watch, but engage now with a conditional term sheet — this is a real workflow tool with genuine field ROI, not a science project. The single strongest reason for: ShelfSense is capturing paying CPG demand ($26k MRR across 12 brands, reps 3x faster) in a sector riding a 36% tailwind, and the technology is de-risked. The single strongest reason against: it sits squarely in the thin-wrapper trap (competitive headroom 41/100) — commodity CV that Trax and retail-media incumbents already own at scale and can bundle for free, with no disclosed data-license rights or LTV/CAC to prove a durable moat. Entry plan: do not wire on current disclosure. Resolve the ~12x MRR discrepancy ($26k vs. $312k parsed) as an R&W condition, demand a held-out field-accuracy benchmark and contractual image-retention/training rights before closing. On clean diligence, lead $1.75M for ~8.8% at ~$16.4M pre, reserve $2.625M for pro-rata, sizing ~2% of the portfolio.

Narrative engine: model anthropic (per-part source not recorded)

Entry strategy

Ticket: $1,750,000 target (range $875,000–$1,750,000)
Target ownership: 8.8%
Valuation band (pre-money): $8,384,000 / $16,440,000 / $32,880,000
Modelled return (not a forecast): 6.89x multiple (16.4x base) · horizon 7yr · loss prob 59%

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 ~2.1% of a diversified venture portfolio (fractional-Kelly, conviction-scaled). Reserve 2,625,000 USD for pro-rata follow-on.

Score breakdown

Market size & growth — 57/100 (weight 14%)
~$70B TAM, 36% CAGR (AI Applications (vertical SaaS)).

Timing / tailwinds — 100/100 (weight 5%)
Sector growth 36% vs. 12% neutral baseline.

Moat / defensibility — 59/100 (weight 16%)
switching costs is the category's mature moat (ceiling 74), but ~61% realized at seed given disclosed traction — an unproven moat is discounted toward the 35 "no demonstrated defensibility" floor.

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

Team / execution signal — 64/100 (weight 28%)
Quantified traction: $312k MRR.

Scientific / tech feasibility — 98/100 (weight 7%)
agentic workflows, domain eval harnesses, retrieval + tool orchestration

Regulatory / legal headroom — 74/100 (weight 7%)
Regulatory intensity 40% (higher = more legal drag).

Competitive headroom — 41/100 (weight 8%)
Competitive intensity 85%. thin wrapper risk — value must accrue above the model layer.

Analyst council

Research Scientist — Shelf CV is technically mature — 98/100 overstates the differentiation; execution and data moat, not feasibility, are the real questions.

+ The 98/100 feasibility score is directionally right on 'can it work' but conflates that with 'is it hard/defensible.' Retail shelf recognition is a solved-enough problem: SKU-level classification and planogram matching run on off-the-shelf detectors (YOLO-class / vision transformers) at 90-95%+ mAP on clean shelf imagery. Trax, Standard AI, and even Google Cloud Vision already do this — so the sci

+ The genuine technical risk is NOT model accuracy in demos but robustness in the field: phone photos vary in angle, glare, occlusion, and store lighting; accuracy on messy real shelves (reflective packaging, near-identical SKU variants, private-label lookalikes) can drop 15-30 points vs. curated benchmarks. The moat, if any, is a proprietary labeled dataset of long-tail CPG SKUs plus a domain eval

+ 'Agentic workflows / retrieval + tool orchestration' in the frontier tags is largely irrelevant hand-waving for this product — the core is supervised computer vision, not LLM agents. If the pitch leans on agentic-AI framing to justify the score, be skeptical: it signals narrative-chasing rather than the actual engineering (edge inference, SKU catalog management, planogram diff logic). Real value-a

+ Traction is real but early and internally inconsistent: description says $26k MRR / 12 brands / 60k shelves; parsed plan metric shows ~$312k (annualized run-rate?). This discrepancy must be reconciled — a 12x gap between stated MRR and parsed revenue is a data-integrity flag even though it's not auto-detected as a RED FLAG.

! Thin-wrapper / commoditization: core CV is replicable in 6-12 months by a competent team on open models, and incumbents (Trax, Standard AI, ITsML) plus retailer-native tools already occupy the space. Without a proprietary SKU dataset moat, gross margins (~70%) compress as this becomes a feature, not a product.

! Field-robustness gap: no disclosed accuracy metrics on real-world uncurated photos. If false-positive out-of-stock/compliance alerts exceed ~5-10%, reps lose trust and churn — the '3x faster' claim is self-reported and unaudited. Demand a held-out field accuracy benchmark before conviction.

! Metric inconsistency ($26k MRR vs. $312k parsed revenue) plus no disclosed LTV/CAC or bottom-up TAM means neither stress test nor TAM triangulation could run — feasibility is de-risked but the business case is not.

Data Analyst — 69/100 unit-econ score is generous on assumption, unproven on data: 70% GM is plausible but CAC/LTV, payback and churn are all blank

+ The 69/100 leans on a sector-reference ~70% gross margin and 35% capital intensity — reasonable for CV SaaS, but there is zero disclosed CAC, LTV, payback or churn, so the score reflects a template not this company. I'd mark it 'insufficient data' rather than 69; the number flatters a blank.

+ Note the MRR discrepancy: description says $26k MRR (12 brands "H $2.2k/brand/mo, ~$312k ARR) while the model's parsed field reads $312k as MRR. If ARR is ~$312k, that is a real seed with land-and-expand potential; if MRR is truly $26k the team score's '$312k MRR' claim is wrong. Confirm before anything else.

+ Unit-econ thesis hinges on seat/shelf expansion within each brand: 60k shelves/mo across 12 brands is shallow penetration of a national CPG's store base (a single big CPG audits 100k+ facings monthly). Net revenue retention >120% would confirm the land-and-expand story; flat or <100% NRR kills it given the low absolute MRR.

+ Structural counter-argument: 'thin wrapper' risk (competitive headroom 41) directly threatens margins — if value sits at the model layer, inference costs compress the 70% GM and a well-capitalized incumbent (Trax, Snap/EDGE, ITL) reprices the category. Durable margin requires proprietary planogram-eval data + workflow lock-in, not just CV.

! No LTV/CAC/payback disclosed and stress test not run — a 3x-faster-visit ROI is a feature claim, not proven willingness-to-pay; if CAC exceeds ~12mo payback at $2.2k/brand ACV the model is uninvestable at seed.

! Established shelf-analytics incumbents (Trax, ITL/Shelfgram, Snap) already offer comparable CV compliance; a $3.5M seed buys little defensibility against their data moats and CPG relationships — churn to a bundled incumbent is the base-rate kill scenario.

! Concentration and TAM unproven: 12 brands, no bottom-up TAM triangulation. The ~$70B sector TAM is not ShelfSense's SAM — realistic SOM is field-audit software spend (low single-digit $B), and loss of 2-3 anchor brands would erase a third of revenue.

Economist — ShelfSense: real workflow ROI but a $70B TAM headline masks a much thinner serviceable retail-execution wedge

+ The 57/100 market score is fair-to-generous. The '~$70B AI vertical SaaS TAM' is a category abstraction, not this company's SAM. Retail-execution/shelf-intelligence software (Trax, Repsly, Wiser, ISR/Pemberton-type tools) is realistically a $2-4B serviceable market; the addressable slice ShelfSense competes for is CPG field-rep audit tooling, not all vertical AI. I'd anchor to that narrower pool,

+ Demand is inelastic where labor is the substitute: '3x faster store visits' converts directly to headcount/route economics for a CPG rep force costing $60-90k/rep fully loaded. That is a hard-dollar ROI pitch with genuine willingness-to-pay, which supports the ~70% gross-margin assumption. This is the strongest economic argument for the deal.

+ Rents accrue to whoever owns the CPG data relationship and planogram integrations, not the CV model. $26k MRR across 12 brands = ~$2.2k/brand/mo — small, and pilot-shaped. Switching cost (the 59/100 moat) only becomes real once ShelfSense holds historical compliance data and is wired into brands' merchandising systems; at 60k shelves/mo that flywheel is unproven.

+ Note the model's internal inconsistency: 'Team' cites $312k MRR while traction says $26k MRR — the parser appears to have annualized. Underwrite the $26k monthly figure; the composite's 64 team score is inflated by this

artifact.
! Thin-wrapper / value-capture (structural, competitive headroom 41/100): commodity CV + LLM tooling means the defensible layer is integrations and proprietary shelf data, not the vision model. Trax and retail-media incumbents already own shelf-scanning at scale and can bundle it free into media/data deals — the disclosed-incumbent-substitution risk is real.
! Concentration and churn: 12 brands, ~$2.2k ARPA. Losing 2-3 pilots halves revenue. No LTV/CAC disclosed (stress test not run) and no bottom-up TAM (triangulation not run), so retention and true SAM are unverified — the two facts that most determine whether this is a business or a feature.
! Macro sensitivity: CPG marketing/field-execution budgets are cyclically cut first; a 3x-efficiency tool can be reframed by buyers as a headcount-reduction lever (shrinking seat count) rather than net-new spend, capping account expansion.

Corporate & Regulatory Lawyer — 74/100 legal headroom is fair-to-generous; ShelfSense is a genuinely low-regulatory AI use-case, but MRR discrepancy and IP silence are the real gates.

+ Agree with 74/100, lean slightly higher on the merits: retail planogram/OOS detection is a benign use-case with no HIPAA/GLBA/hiring exposure. The main federal surface is FTC Act §5 on AI accuracy claims — '3x faster' and 'compliance' marketing must be substantiated or it becomes a deception theory. Cabin claims to measured pilot data.
+ Data/privacy exposure is thinner than typical AI SaaS: shelf photos are property/product imagery, not PII, so CCPA/CPRA and the state patchwork are largely dormant unless reps' faces/bystanders are captured — require an incidental-face redaction policy and a B2B DPA passing store/CPG-owner image rights downstream. Colorado AI Act (effective 2026) targets 'consequential decisions' and does not reac
+ IP posture is undocumented and this is the actual gap the memo must flag, not the regulatory score: no patents, trade-secret, or model-ownership disclosure. Given the composite's own 'thin wrapper' structural risk, verify ShelfSense owns its training data/annotations and eval harness rather than relying on a third-party vision API whose ToS could reclaim or throttle it. Value must accrue in propri
+ Round mechanics are clean and standard: $3.5M seed via Reg D 506(b) (no general solicitation) with Form D in 15 days, or 506(c) with accredited verification if they market publicly — confirm every investor's Rule 501 status matches the exemption. Insist on 1x non-participating pref, broad-based weighted-average anti-dilution, pro-rata and information rights, and a board observer seat.
! Data-flow rights, not privacy law, are the sharp legal risk: if CPG customer contracts don't grant ShelfSense a license to retain/aggregate shelf images for model training, the proprietary-data moat is legally hollow and the 'thin wrapper' penalty (competitive headroom 41/100) hardens — the company owns software but not defensible data.
! Note the $26k MRR (description) vs. $312k MRR (parsed model input) contradiction — a ~12x discrepancy the model treated as fact. This is a diligence/representation-warranty issue: financial reps in the SPA/SAFE side letter must be verified, and any traction misstatement to accredited investors is itself §5 / anti-fraud exposure.
! FTC §5 substantiation: if '3x faster visits' or 'planogram compliance' underperforms in the field, a CPG customer or the FTC could frame marketing as deceptive; require claims to be tied to auditable eval-harness metrics and add a customer-facing accuracy disclaimer.

Market data sources

- Global Market Insights (2026) — Generative AI $83.3B in 2026 !' $988.4B by 2035 at 31.6% CAGR
  https://www.gminsights.com/industry-analysis/generative-ai-market
- Fortune Business Insights (2026) — Generative AI $161B in 2026 !' $1,260.2B by 2034 at 39.6% CAGR
  https://www.fortunebusinessinsights.com/generative-ai-market-107837

Assumptions & limitations

- Market size / growth for AI Applications (vertical SaaS) is anchored to Global Market Insights (2026): Generative AI $83.3B in 2026 !' $988.4B by 2035 at 31.6% CAGR. Full citations are listed under "Market data sources".
- Signal coverage: ~44% of the score is backed by the plan's own disclosed metrics (1 quantified field); the remainder uses AI Applications (vertical SaaS) sector priors — add financials to raise it.
- 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.

This memo is generated by an AI screening tool for research purposes and is not investment advice, an offer, or a solicitation. Figures are model estimates, not guarantees.