

AMI Labs: Independent DD Research Report

Independent market and moat analysis for investment, partnership, or category entry decisions. Prepared for a Series A investor evaluating AMI Labs as a lead investment opportunity. Research period: Q1 2026. This report covers technology moat, market sizing, team assessment, competitive exposure, and investment verdict.

DUE DILIGENCE SCORECARD

DIMENSION	SCORE	RATING	ANALYST NOTE
Technology Moat (JEPA Architecture)	9/10	■■■■■■■■■	V-JEPA 2 zero-shot robotics result is peer-reviewed and reproducible. No competitor has matched latent-space efficiency at scale.
Founding Team Quality	10/10	■■■■■■■■■	LeCun (Turing Award), Xie (DeepMind/Meta), Rabbat (FAIR Montreal). Highest-credentialed founding team in the physical AI category.
Market Timing	8/10	■■■■■■■■■	LLM ceiling is becoming operationally visible in enterprise. Industrial and clinical buyers are actively evaluating alternatives.
Competitive Moat vs. Hyperscalers	6/10	■■■■■■■■■	Google DeepMind and OpenAI have 10–100x more compute. Moat depends on architectural superiority holding — a testable thesis.
Commercial Traction	5/10	■■■■■■■■■	Nabla partnership is credible but pre-revenue. No signed enterprise pilots disclosed as of Q1 2026. Traction is the key open question.
Go-To-Market Clarity	6/10	■■■■■■■■■	Clinical AI path via Nabla is clear. Industrial robotics path requires data partnership deals that are slow to close.
Capital Efficiency	7/10	■■■■■■■■■	\$1.03B seed is large but team is deliberately small (~12). Burn rate is research-intensive, not sales-driven. Runway is strong.
Regulatory / Geopolitical Risk	5/10	■■■■■■■■■	EU AI Act creates compliance overhead. FDA/CE clinical pathway adds 12–24 months to healthcare deployment. Not fatal but real.
Open-Source Strategy	8/10	■■■■■■■■■	'Linux of World Models' narrative is differentiated and talent-attracting. Open weights plan aligns with Nvidia and academic ecosystem.
Exit / Return Potential	9/10	■■■■■■■■■	If physical AI thesis holds, AMI is a generational company. \$3.5B pre-money vs \$50T TAM implies massive upside at Series A.

KEY DUE DILIGENCE QUESTIONS & FINDINGS

Is the JEPA architecture genuinely defensible, or can hyperscalers replicate it with scale?

DEFENSIBLE

FINDING: JEPA is not a secret — the research is published. The moat is execution, not secrecy. LeCun's team has a 3–5 year head start in V-JEPA iterations, proprietary training datasets from Nabla and industrial partners, and the talent density to iterate faster than a hyperscaler team running 20 competing projects. Replication risk is real but underestimated by bears — the value is in the trained model weights and domain-specific data, not the architecture paper.

What is the realistic path to revenue, and when?

MODERATE RISK

FINDING: The clearest near-term path is Clinical AI via Nabla — Nabla has existing hospital relationships, CE marking experience, and a co-founder overlap with AMI's CEO. A formal pilot contract in 3–5 EU/Singapore hospital systems is achievable within 12 months. Industrial robotics deals will take longer (18–30 months) due to data partnership cycles. Investors should expect a pre-revenue Series A with pilot LOIs as the key proof point, not signed ARR.

Is the \$3.5B pre-money valuation justifiable at seed stage?

PRICED FOR
PERFECTION

FINDING: By conventional seed metrics, the valuation is aggressive — there is no revenue, no deployed product, and ~12 employees. However, comparables support it: Mistral AI raised at a \$6B valuation pre-product; Physical Intelligence raised at a comparable stage. The LeCun name functions as a category-creation asset. The key risk is that the \$3.5B becomes an anchor that constrains Series A pricing if traction is slower than expected.

What does the competitive landscape look like in 24 months?

WATCH CLOSELY

FINDING: The most dangerous scenario is Physical Intelligence (Pi) expanding into EU industrial markets before AMI secures anchor partnerships. Pi is better-capitalised at the robotics execution layer, and their $\pi 0.2$ model is already deployed in 68 manipulation tasks. AMI's counter is the open-source strategy: if V-JEPA 2 open weights create an ecosystem of fine-tuned vertical models (similar to Llama's effect on LLMs), Pi's closed-model approach becomes a liability. This is a race AMI can win — but speed matters.

INVESTMENT VERDICT

OVERALL RATING

**STRONG
BUY**

Weighted DD score: 7.3 /
10

AMI Labs is the most credible architectural bet against the LLM paradigm currently available to institutional investors. The founding team is unreplicable, the technology thesis is peer-validated, and the market timing is right. The primary risks — commercial traction and valuation anchoring — are knowable and manageable within a 12-month pilot window. Recommended entry: lead or co-lead Series A at a valuation cap of \$8–10B, contingent on 3+ signed pilot LOIs in clinical or industrial verticals.

PRICING & ENGAGEMENT

Starts at USD 5,000

Flat-fee engagement. Deliverables: full DD scorecard, key question analysis, competitive exposure assessment, and investment verdict memo with entry recommendation. Delivered in 10–14 business days.

techwhoknows.com · jin@techwhoknows.com ·
TechWhoKnows

© 2026 All rights reserved. Sample brief — illustrative only.