

AI's New Moat Is the Work Loop

VC Tech Radar

2026-08-13

AI's New Moat Is the Work Loop

By VC Tech Radar • August 13, 2026

Tracks the strongest signals around self-improving domain labs, verifiable agent environments, robotics reliability, post-training economics, and outcome-based vertical AI.

1. Funding & Deals

Event Horizon Labs is an early-stage bet on self-improving domain expertise. Dalton Caldwell says he led the Series A in EHL, a new quantitative research lab founded by Owen. Its thesis is to automate the type of work Owen previously did at Citadel with self-improving AI; Caldwell calls Owen one of YC's most talented AI researchers and says several teams are pursuing the same direction. [1]

2. Emerging Teams

Suhail's autonomous-AI-scientist project is moving from experiment to compute-backed team. The build log says the project has a validated basic RLVR post-training stack, made its first hire while recruiting for post-training or low-level optimization, and grew from one person to three. After a key research component worked but needed scaling—and GPUs were delayed by networking issues—Suhail reported that much greater quantities of compute were locked down and ready. [2, 3, 4, 5, 6, 7]

The diligence question is whether the validated post-training loop compounds once the new infrastructure is online; the founder's own updates frame scaling research and compute access as the immediate bottleneck. [6, 7]

3. AI & Tech Breakthroughs

Agent training is becoming an environment-and-verifier business. A current industry talk describes the shift from low-skilled crowdsourced behavior-

cloning data toward expert-built environments containing realistic “worlds,” high-fidelity app clones, and tasks with rubric or unit-test verifiers. The speaker reports 2.5 million expert hours in the second quarter; in one 1,800-task post-training run using about \$500,000 of compute, the reported score rose from 4.7% to 26% and generalized to other benchmarks. [8]

That points to a valuable infrastructure layer beyond raw model supply: humans still have to measure performance beyond a model’s frontier, while the next data wave is moving toward 100–1,000-hour tasks and social interaction—an area where the speaker says only about 1% of evaluations currently measure performance. [8]

Robotics reliability is being attacked with reinforcement learning and memory, not demos alone. The robotics talk reports a general-purpose value function trained on robot experience, with human intervention used to avoid dead-end trajectories; the resulting policy ran for 13 hours, exceeded 90% espresso success, and gained roughly 2x box-building throughput from the RL stage. [9] Multi-timescale memory—short video memory plus compressed text for longer history—enabled a non-repetitive kitchen-cleaning task lasting 10–15 minutes, while a single Pi07 model reportedly matched or outperformed fine-tuned specialists and generalized to scarcely represented appliances and a new robot platform. [9]

4. Market Signals

The near-term moat is shifting from the base model to post-training, domain data, and cost control. The post-training thesis is to use frontier models to reach product-market fit, then use product data to encode a company’s distinctive taste and expertise. Because applications can be cloned from screenshots, the tuned model becomes the proposed defensible asset; the speaker says post-training can also reduce serving cost by 5–10x. [10] The same talk warns that reward hacking and training-to-serving drift can make apparent progress useless, making repeatable evals and production A/B tests part of the product rather than a launch afterthought. [10]

Vertical agents are being framed around cash outcomes rather than task volume. Stuart’s collections analysis covers billions of dollars of receivables and finds that the largest 10% of past-due invoices hold 65.7% of overdue dollars. In its data, 81.7% of outbound collection emails require no human involvement, while escalations, disputes, manual calls, and broken-promise follow-ups remain human-resolved; agent-run teams average 1.35 outbound asks per \$1,000 collected, with three in five completed tasks resolving fully automatically. [11] The product pattern is clear: automate the search and repetitive asking, while reserving people for relationships and exceptions.


AI is changing founder leverage faster than it is changing the fundamentals of company building. Garry Tan argues that agentic and vibe coding can make one person “400 of that person” and expects a wave of experi-

enced 35–45-year-old technical founders. Paul Graham’s counterweight is that almost all standard startup advice still holds: the core remains building what users need and finding growth. [12, 13, 14, 15] The investment implication is to widen founder sourcing toward experienced operators without relaxing scrutiny on distribution and durable demand.

Local and open inference is becoming a meaningful deployment channel. Hugging Face says Transformers.js crossed 10 million monthly downloads—nearly 10x its level six months earlier—and attributes local adoption to free, private execution amid compute shortages and cyber-attack risk. Cohere’s North Micro Vision adds to the supply side: its smallest vision-language model is open-source under Apache 2.0 for document understanding. [16, 17]


5. Worth Your Time

- **Watch RL Environments Explained: How AI Agents Learn Real-World Work.** The clearest explanation in the corpus of why agent progress depends on realistic apps, expert-built tasks, and verifiers—not just more model-generated trajectories. [8]



RL Environments Explained: How AI Agents Learn Real-World Work / Brendan Foody, Mercor (2:59)


- **Watch Chelsea Finn: This is the State of the Art in Robotics.** Focus on the memory and generalist-policy sections for a concrete view of how robotics is moving from isolated skills toward longer, compositional tasks.



[9]

Chelsea Finn: This is the State of the Art in Robotics (17:48)

- **Watch Post-Training Is How You Keep Your Taste.** Useful for deciding when a startup should rent frontier intelligence and when its product data, evaluation loop, and economics justify owning a specialized



model. [10]

Post-Training Is How You Keep Your Taste / Fireworks CEO Lin Qiao

(25:45)

- **Read Paul Graham on startup advice in the AI era.** A short corrective to claims that AI has already overturned the basic work of finding a product users want. [14, 15]
-

Sources

1. X post by @daltonc
2. X post by @Suhail
3. X post by @Suhail
4. X post by @Suhail
5. X post by @Suhail
6. X post by @Suhail
7. X post by @Suhail
8. RL Environments Explained: How AI Agents Learn Real-World Work | Brendan Foody, Mercor
9. Chelsea Finn: This is the State of the Art in Robotics
10. Post-Training Is How You Keep Your Taste | Fireworks CEO Lin Qiao
11. X article by @realtarek
12. X post by @a16z
13. X post by @a16z
14. X post by @paulg
15. X post by @paulg
16. X post by @ClementDelangue
17. X post by @cohere