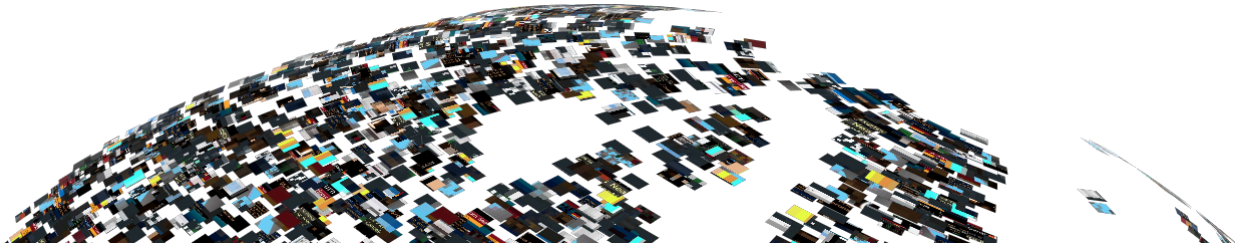
ALEXANDRIA

Written by hand by Raphael Volpert and Nils Anton Kratz



Alexandria is the neolab transforming enterprises into autonomous organizations.

Enterprises are agentic swarms in human form.
Employees act, receive feedback, and adjust.
This is reinforcement learning in biological neurons, not weights.

Shifting execution from humans to agents demands models that understand organizational reward structures.

Quant firms already show us how autonomous enterprises will look like.
At Jane Street, every trade has a price and a measurable outcome.
Those outcomes feed back into internal accumulated knowledge in form of research and models, which shape the next trading decision.
Actions are taken autonomously because organizational reward signals are explicit and quantifiable.

In all other enterprises, feedback loops are ambiguous.
Organizational complexity decouples actions from their true consequences:

- a part an engineer designs in CAD that fails inspection in the MES a week later
- a pricing model an analyst builds that the sales team quietly works around
- a workflow an FDE configures that the customer rolls back the next month

Humans can infer feedback from tone, delays, and scattered systems.
For models these unstructured signals are random noise.

As long as agents cannot learn from organizational feedback, enterprises cannot run autonomously.

Alexandria gives all organizations the feedback loop of a quant firm.

To do this, we:

1. Build human datalabs that privately capture enterprise expertise.
2. Turn the whole enterprise into one continually learning RL environment.
3. Run agentic swarms towards measured outcomes until the company runs autonomously.

Co-Founders

Raphael Volpert: CEO. Background in CS and physics. Sold his first company to a YC startup at twenty. Received Hummingbird VC's Magnificent Grant to benchmark the effects of nootropic compounds using brain organoids. Founder of Calculus House SF where he built robotic wetlabs and researched fUS-brain computer interfaces. a16z scout.

Nils Anton Kratz: CTO. Background in CS and mathematics. CTO of an $8M ARR company at sixteen. Ranked first in statewide ML research competitions. Hired by BMW through its talent program while still in school. Member of the Core Architecture Team for BMW's 40+ PB Cloud Data Hub.

Angels

Alexandria is the most sensitive enterprise deployment you can think of.
Our angels give us unique insight into enterprise sales in IP-heavy, regulated industries:

- **Dr. Robert Fink**, inventor of Palantir Foundry, CTO of Helsing
- **Andrej Henkler**, earliest investor in Palantir, GP at Leblon Capital
- **Saturnin Pugnet**, who founded Worldcoin with Sam Altman
- **Sebastian Enderlein**, CTO of DeepL, prev senior manager at Uber and Salesforce
- **Jonas Rieke**, co-founder and COO of Personio
- **Niklas Jansen**, GP of Interface Capital
- **Philipp Moehring and Christian Eggert** of Tiny VC
- **Noor Sweid**, founder and GP of Global Ventures
- **Maximilian Eichler** of Isembard
- **VPs of Sales** at MongoDB

Core Thesis

Coding agents became autonomous because code is explicit and recorded.
They are trained on a decade of granular GitHub commits and purpose-built RL environments.
Every commit is a step, every merged or rejected pull request a reward.

Most knowledge work in enterprises is tacit and unrecorded.
Companies run on the social judgement and domain intuition of their experts.
All of it happens on desktops. Every day unrecorded, this asset is left behind.

Alexandria builds enterprises their own human datalabs:

Ambient

Ambient is a personal world model, which continually learns from your desktop activity.

It uses state-of-the-art vision language models to understand your work across any application, from CAD and Excel to internal tooling or biotech software.
All processing is powered by Alexandria Confidential Cloud, combining the privacy of local models with the intelligence of cloud inference.

Returns are immediate:

1. Ambient identifies high ROI tasks and launches long running agents in your existing agentic stack (MS Copilot, Claude Cowork, ChatGPT work, Grok Bot or OpenCode).
2. Now every agent runs with the full context of your work.

Multiplayer

A company's expertise is spread across thousands of desktops.
Multiplayer brings every Ambient together, in one continually learned model of the company.
Built for agents to query, it provides the most valuable dataset of your enterprise: its full operational context.
Alexandria's RLCD-based privacy classifier ensures that anything shared to your organization's Multiplayer is fully PII-redacted.

Continual Learning and Parametric Memory

Human data labs create a flywheel.
The longer Ambient observes a user, the better it finds signal in the next desktop activity timeframe.
This is why we build on continual learning and parametric memory:

Continual learning is moving from high-sample post-training to one-shot updates in weight space.
Keeping a company's knowledge in a store and querying it at runtime is retrieval, and retrieval is not learning.
The last months showed parametric memory works: separated from reasoning (DeepSeek), updated without forgetting (Meta), written into weights with no training run (MIT).

The Bitter Lesson vs Enterprise Incentives

Solely from bitter lesson standpoint, separating enterprise world models (Multiplayer) from frontier intelligence seems suboptimal:
Scaling laws point to one organization and one model that absorbs the IP of the whole world.
However that view breaks on enterprise incentives:

As long as different organizations compete, their world model is their one advantage, that cannot be rented.

With frontier intelligence providers changing every few weeks:
Companies will rent capability (SOTA tokens) and own their world model.

A model trained just on trajectories learns to imitate but never to drive an outcome:
raw trajectories contain sequences of actions and leave their reward signal unlinked.

Labs now pay Mercor and Scale to rebuild that reward signal from scratch, while every company already holds the full data.

In most knowledge work the reward is sparse and delayed: it may land weeks later, on another desk. Alexandria ties it back to the trajectory that earned it.

With every trajectory carrying its reward, long-horizon credit assignment becomes possible:
value functions trained on those rewards trace each outcome back to the steps that drove it, and score the next action towards a target outcome.

The whole company becomes a continually evolving reinforcement learning environment:
Outcomes for its agents to work against, verifiers that check it while it runs, and evals that measure any model on the company's real work.
A company now owns the feedback loop a quant firm has for every trade.

Autonomous Orgs

Alexandria now shifts enterprise execution towards agentic workforces:
swarms of agents that receive an outcome and optimise towards it.
Swarms act, receive feedback, and adjust.

This is the metamorphosis Alexandria puts a company through.
First, human trajectories become long-running clones of the experts who produced them.
Then, by treating the whole enterprise as one reinforcement learning environment, the clones become swarms that surpass the experts.

As memory and harness move into the model's own weights, the company becomes a self-improving autonomous organization.

Traction

Three months since inception:

- Deployed at frontier research labs and enterprises, among them a stealth robotic wet lab in San Francisco, a European HR software unicorn and a precision manufacturing startup
- In security and legal review at nine enterprises, among them a German premium carmaker, a global pharma company, a Big Four firm and a European defence AI company
- Scoping solutions with over 50 more companies across banking, insurance, pharma, automotive, industrials and consumer goods

We sell top-down to the executive accountable for the AI transformation.
We start bottom-up: Ambient on one workflow, with security, legal and the works council involved early.

Ad Astra.

Alexandria emerges from stealth Watch on LinkedIn ↗