

Your agent stack burns API credits on a task a 30-million-parameter classifier could own. Every routing call, every rerank, every moderation pass goes to a frontier model at frontier prices, the latency compounds, and the bill climbs. Then the "should we fine-tune or just add RAG?" debate eats another sprint and nothing ships.

The standard advice is to learn ML properly first: take the course, derive the gradients, come back when you can train a network from scratch. It fails because you are not publishing a paper; you are putting one small, owned model into production beside an LLM and keeping it alive.

ML for Agent Builders is written for that gap. It hands you BUILD (a loop, not a lab project): target the metric before the architecture, train and eval like product, and inspect drift before it degrades the system. The frontier model orchestrates; your narrow models execute, and BUILD makes them earn their place.

You will learn to:

- Decide build vs buy vs prompt with one matrix, not three debates per sprint.
- Pick fine-tuning or RAG by capability and cost, and know when to combine them.
- Generate synthetic data fast and verify it faster, with guardrails that catch failure modes.
- Stand up an eval harness that runs a regression suite, with real ship gates, before every deploy.
- Set a drift budget that tells you when to retrain, when to roll back, and when to do nothing.

This is not data science for spectators or MLOps for a fifty-person platform team. It is operator-grade ML: protocols you run with a model in production, a latency budget that bites, and an agent on hand to help.

The result is concrete. You stop renting frontier compute for jobs a small owned model does cheaper, faster, and under your control, and you ship models you can defend, monitor, and sunset on purpose.


If you build production systems and you are done paying retail for inference you could own, start here.

Len P. van der Hof, MSc

He builds operating systems for the strategic self. Treats entrepreneurship, AI and machine learning, marketing, philosophy, psychology, and health optimization as one engineering problem. MSc in Strategic Entrepreneurship, Rotterdam School of Management, Erasmus University.

ML for Agent Builders

Auto-training, fine-tuning, and eval loops for LLM-powered systems



Len P. van der Hof

Systems for the Strategic Self