



Trajectory

Own Your Intelligence

Trajectory turns real product usage into AI that continuously improves.



AI is intelligent. It isn't experienced.

Off-the-shelf models don't know your users, your workflows, or what good looks like in your product.

01 Expensive

Frontier tokens costs are scaling faster and killing margins

02 Unreliable

Agents break on the long-tail tasks your users actually hit.

03 Frozen

What your product learns in production never makes it back into the model.

04 Thrown away

Every interaction leaves a signal - edits, retries, corrections. Today they're discarded.

A frozen model has intelligence. A learning model has experience - and experience compounds.



Trajectory - The Platform for Continual Learning

Two easy integration points:

01 • THE CALL

Every request and response

Wrap the client you already use, and messages, tool calls, token counts, and latency stream in.

```
llm = Trajectory(OpenAI(api_key))
```


02 • THE OUTCOME

What you record as good/bad

All the accepts, edits, retries, and escalations that occur in your product, or a clear definition of success per task.





THEN THE LOOP RUNS, CONTINUOUSLY



03 - THE OPTIMIZATION

The Control Plane

An interface to turn production traces into post-trained models, that you can eval and deploy



Harvey × Trajectory

Opus-level legal work on a model Harvey owns.

Opus-level performance




83% on Harvey LAB vs 84% for Opus 4.6 - share of rubric criteria passed on 1,200+ end-to-end legal tasks.

10× cheaper than Sonnet, 3× faster

At a fraction of the cost per token, on open weights Harvey owns and can deploy anywhere.

Under 24 hours


From base model to post-trained, evaluated model - delivered end to end.



 Rogo × Trajectory

Most of the way to Sonnet,
on a model they own.


24% → 36% on BigFinanceBench · 10× cheaper than Sonnet, 3× faster ·
output halved: 10.5K → 4.9K tokens per question




 Decagon × Trajectory

A 35B open model, ahead of the frontier.




Joint exact match 86 vs Opus 4.6 at 83 · severity 90 vs 84 · on
hallucinated cases with customer overrides, at open-model cost




Better on your metric, at open-model cost - and it keeps improving.



Working together

01 POC design	02 Train & eval	03 Deploy & compound	04 What we measure
 Together, we define the use case, success criteria (accuracy, cost, latency) and intake mechanisms.	 Ingest data by wrapping model calls via SDK and providing reward. Train an open source model, and eval against other models	 Deploy the model, collect more production data, repeat.	 Outcomes in production, against your baseline: performance up, cost down, latency down.



Built by research and industry experts



Ronak Malde (CEO)

Trained SWE-1 at Windsurf. Led to the \$2.4B DeepMind acquisition.



Michael Elabd (CTO)

Built continual learning infra for robotics. Reasoning co-lead, Gemini Robotics 1.5.



Arjun Karanam (CPO)


Trained tracking AI on Vision Pro. RL for multimodal foundation models.




We are a group of creative technologists building continual learning.

We turn every interaction with your product, every correction, retry, and edit into signal that grows the intelligence inside it.


We have gathered a stellar team from



& work with



Our supporters include:





Intelligence

in motion

Trajectory

Every great AI product should own its intelligence.

Companies rent intelligence today. Trajectory is the infrastructure for owning it:
your product → your data → your learning loop → your model → your advantage.

Questions?

sales@trajectory.ai

