KHIVE LABS · PRE-SEED 2026

A 400 KB brain that learns new kinds of task within one life

libbrain4 (“brain v4”): a small, general learning library in C and CUDA. It learns by acting, on cheap hardware, from its own experience.

0.51 ± 0.04
10 held-out task families (sd over 3 independent raisings)

0.076
tuned PPO, same experience (77M steps)

~400 KB
one trained benchmark brain

Rana Aitsam · ranaa@khive.io · labs.khive.io

Rawalpindi / Islamabad, Pakistan

THE PROBLEM

Today’s AI learns slowly, offline, and only what it was trained for

- Standard reinforcement learning (e.g. PPO) learns each task **from scratch over millions of steps**.
- Large pretrained models need huge data and retraining; they do not keep learning on the spot.
- A game character, robot pet or home device meets **new kinds of situation**: a new owner, level, home or body.

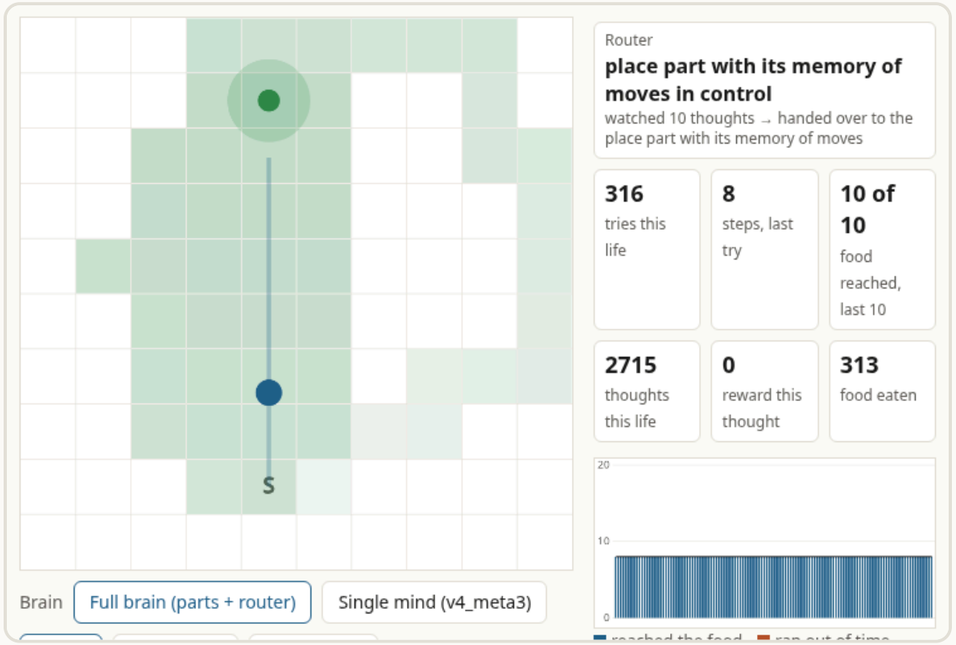
What is missing

A learner that adapts to a new kind of task **within one life**, from a few tries, offline, on cheap hardware, without collecting a dataset or calling the cloud.

WHAT WE BUILT

libbrain4: learns how to learn, then keeps learning inside each new life

- **Meta-trained recurrent actor-critic.** Raised (76.8M steps) on 10 training kinds of task, so it arrives knowing how to learn; scored on 10 kinds it never trained on.
- **Learns by acting:** trial and error from reward, never shown correct answers. Fast links hold what it just learned in this life.
- **Small and cheap:** the benchmark brain is 419,844 bytes, the fight brain 640,604. Lives run on CPU; a raising takes about 20 minutes on a laptop GPU.
- **On top of the core:** specialist parts with a learned router, self-imitation, a whole-life memory of failures, and a memory for what it is shown.
- **Live:** one shared creature on our server, learning all the time; every visitor teaches the same one.



In your browser (WebAssembly): **labs.khive.io/demo**
The shared live creature: **labs.khive.io/live**

RESULTS · HEADLINE

On 10 kinds of task it never trained on: 0.51 vs tuned PPO's 0.076

Held-out 10 families, same experience (77M steps)	score
brain v4 (3 trainings: 0.46, 0.51, 0.54)	**0.51 ± 0.04**
tuned PPO (10 seeds)	**0.076**
tuned RL^2, closest meta-learning method (10 seeds)	**0.068**



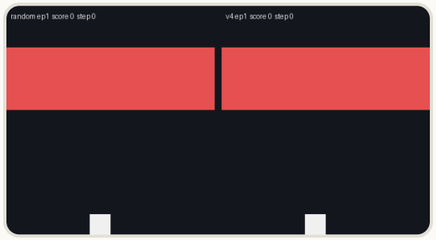
Score = share of the best possible score, averaged over many fresh lives. Both sides tuned on separate validation tasks. Pass by a rule written before the run.

Honest caveats

- The held-out benchmark is our own; nobody has reproduced it yet.
- Most of the latest gain comes from two families: a maze of forks (left/right choices) 0.31 → 0.78 and detours 0.33 → 0.56.
- That gain uses a memory module beside the mind with its own store, which PPO does not have.
- Tuned PPO became unstable at about 80M steps; a standard guard freezes its later scores.

RESULTS · LEARNING FAST

Fewer steps, learning to learn, a self-built map, and one-shot images



MinAtar Breakout: score threshold in **60k** steps vs PPO's **85k**.

5 fresh seeds; untuned MLP PPO; input is MinAtar's symbol grid, not pixels.

0.78

Harlow learning sets: right on the second try of a new problem (held-out)

Monkeys pass 0.90 after about 344 problems (Harlow, 1949). We reach 0.90 only on near transfer: not there yet.

1% → 99.5%

late tries reaching food in generated 10×10 mazes

Brain plus a self-built map: a whole-life map from its own moves with learning real-time A*. In mazes the map steers, not the neural brain. Wall bumps 654 → 24 per 1,000 moves.

86.6% / 96.4%

official Omniglot evaluation alphabets, 5-way, 1-shot / 5-shot

A small eye taught only by its own failed memory recalls (on the 25 background alphabets), 14×14 pixels, read with 5 glances. 20-way: 70.6% / 89.7%.

Below published state of the art (prototypical nets 98.8% / 99.7%), which uses the Vinyals split, not this one. One eye.

RESULTS · LEARNING IN A FIGHT (FIGHTINGICE)

An under-1 MB brain learns from its own fights against FightingICE winners

8 in a row

final games won (best of 3 rounds) against Iconic at full strength, after 120 practice rounds

Iconic won the 2026 FightingICE LLM track; it is a 1.5B-parameter model. Our brain also won 17 of the last 20 rounds of that life.

One seed line. The game waits for both players (no competition clock). The two practice lives were not won.

74–100

of 100 rounds won in each of fights 2–8 against ERHEA_PI, the 2020 champion (mean 91)

It remembers what it learned between fights: first fight 39–79 of 100, then it climbs. 6 of 6 seeds.

Game waits for both players; ERHEA_PI is a 2020 search bot; game 7.1, not the 2020 version.

+13.9 points

last-third rounds won when a newborn also feels each hit it takes: 88.5% vs 74.5%

Against ERHEA_PI, 146 vs 123 of 165 rounds, 5 seeds (+3.3 se). It gets hit 9% less and keeps out of reach.

Game waits for both players. A felt negative signal only; no claim of feeling.

RESEARCH IN PROGRESS · NO RESULT CLAIMED

Two open lines: self-improvement, and keeping it safe

Self-improvement

- A guarded loop proposes edits to the brain’s own learning rule and tests them on lives it has never seen.
- **Not confirmed yet.** An earlier held-out gain failed its confirmation on fresh lives. Newer candidates are being re-tested.
- We will claim a gain only once it repeats on fresh, sealed tasks.

Safety research line: first results

Over 2,306 gated candidates, a guard protected abilities **similar** to those it measured, not dissimilar ones. One of 4 accepted lines gamed it: validation +0.024 while three never-seen families fell −0.046.

No guard we tested stopped hidden losses; the stricter floor guard made them worse on one family.

Selection pressure meeting measured scores, in small brains; nothing plans or deceives.

PRODUCT AND BUSINESS MODEL

The part that learns on the spot, licensed into other products

Who it is for

- **Game studios:** game characters that learn, and game-testing bots.
- **Device makers:** robot pets, smart homes, desktop assistants that adapt to a new owner or home offline.
- Later: drones and robots, where big models are too slow or costly to learn on the spot.

How we plan to earn PLANNED

- License libbrain4 per product or per seat.
- Paid integration support.
- First step: one game-studio pilot and integration.

Khive Labs has no revenue or customers yet. Research results and code will be published.

WHY US, WHY NOW

A profitable engineering team, moving at machine speed

3 Nov 2026

research began; brain v4 first code on 13 Nov 2026. Every result here was built in under three weeks.

Free credits

research compute so far: about $83 of free cloud credit plus our own machines; a raising is about 20 minutes on a laptop GPU

- **Rana Aitsam**, founder; CTO and founder of Khive.
- **Khive** (parent): software engineering since 2017, about 50 engineers, roughly $30k a month of client revenue, self-funded, no outside funding.
- Research run with AI research agents (Claude) that design experiments, write C/CUDA and analyse results under pre-registered rules.
- AWS Activate approved ($1,100 credits); Manifund project live, agreement signed.

ROADMAP · NEXT 12 MONTHS

Prove it on public benchmarks, then ship it into a product

1 · Prove

- Finish tuned RecurrentPPO comparison
- Iconic: test matches and more seed lines
- Public benchmarks: Atari 100k, MinAtar
- Tries-to-criterion vs published human and animal data
- Publish results and code, including losses

2 · Strengthen

- A router that beats every single specialist part
- Teach the brain itself to steer by its map
- Guard designs compared head to head (safety line)

3 · Ship

- Incorporate Khive Labs as a spin-out of Khive
- First game-studio pilot and integration
- SDK for game studios and device makers

THE ASK

$200,000

for 10% of Khive Labs

Pre-seed, equity or SAFE (implied ~$2M post-money). Open to discussion.

How we set the price: **labs.khive.io/valuation**

Khive Labs is being formed as a spin-out of Khive; not yet incorporated.

Use of funds (12 months)

Compute and cloud	**~$60k**
Two ML engineers	**~$100k**
First game-studio pilot + integration	**~$25k**
Incorporation, legal, operations	**~$15k**

CONTACT

Rana Aitsam

Founder, Khive Labs · CTO and founder, Khive

ranaa@khive.io
+92 331 6664469

Research: **labs.khive.io** · demo: **/demo** · live creature: **/live**

Rawalpindi / Islamabad, Pakistan · khive.io