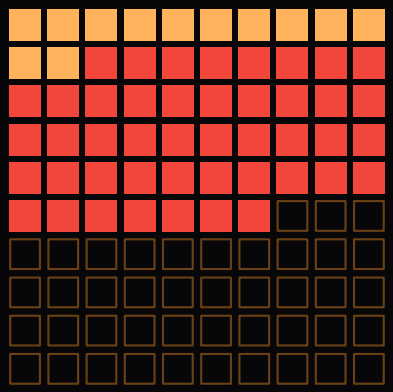
Code that survives an attacker.

We train AI to write code that hackers can't break into.

SAMIR ABIS, FOUNDER · SAMIR@RHEONO.DEV

rheono.

AI writes the code now. Most of what works can still be broken into.



SUSVIBES STUDY · AN AI CODER, PER 100 TASKS

57%
WORKED

11.8%
WORKED AND SECURE

45 in 100 passed every test and could still be broken into.

SECURE · WORKED, NOT SECURE · DIDN'T WORK

Source: SusVibes, arXiv 2512.03262 · SWE-Agent on Claude 4 Sonnet

The AI earns its reward only if the code works and our attacker can’t get in.



Coding companies now train their own AI, and spend big on its practice.

CURSOR · COMPOSER 1.5

20x

more computing power for reward-based practice than Cursor's model before it.

ONE AI LAB · ONE YEAR

$1B+

discussed for practice setups, the tasks an AI trains on.

THE WINDOW They are choosing how to grade their AI today.

Sources: Cursor research blog, 2026 · The Information via TechCrunch, Sep 2025

Our attacker knows what no AI was trained on.

01 KNOW-HOW

15 YEARS

on both sides of security, including closed online markets that attackers use. It isn't on the open web, so no AI learned it.

02 PROOF

35 HOLES

in live apps for 13 bug-bounty programs, 31 by our AI alone, past real defenses.

03 FIRST RESULT

17% → 20.5%

secure, working code after training against our attacker. p = 0.004, on a public benchmark.

THE BIG LABS Our main competitor, and they already buy practice setups from outside.

35 holes reported to 13 bug-bounty programs since 26 Aug 2026 · 31 found by the AI alone

We're raising $3M for 18 months.

TEAM

4 PEOPLE

Me plus three hires, first an engineer who has trained coding AI.

TRAINING

5–10 RUNS

To take our first result to bigger models.

LEADERBOARD

PUBLIC

Our attacker breaks into apps built by each lab's AI, and we rank them in public.

TODAY Contacts at Google, Cursor, Mistral and Qwen. Part of Google for Startups. No commitments yet.

Code that survives an attacker.

SAMIR ABIS, FOUNDER · SAMIR@RHEONO.DEV

rheono.

Answers to the questions we expect.

WON'T THE BIG LABS BUILD THIS?	They are the main competitor, and they also buy practice setups from outside. The hard part is an attacker that knows how real break-ins happen. We sell that.
WHY YOU?	I built the attacker alone, and it finds real holes: 35 for 13 bug-bounty programs, 31 by the AI. Behind it are fifteen years on both sides of security.
WHEN DOES A BREAK-IN COUNT?	Only when 3 checks agree: one AI finds it, a second AI fails to disprove it, and a plain script with no AI repeats it.
WHAT IS PROVEN, AND WHAT ISN'T?	Both halves. It finds real holes in live apps, and training against it raised secure, working code from 17.0% to 20.5% (3 seeds, p = 0.004, preregistered). Next is bigger models.
WHO HAS COMMITTED?	No one yet. We have contacts at Google, Cursor, Mistral and Qwen, and a VC is interested in leading this round. Nothing is signed.