

GINOFIT · IN-HOUSE PRODUCT

Shanghai

A marathon training monitor that grades every workout with AI, rewrites the plan in real time, argues its verdicts in chat, and rewards the discipline to rest. On the desktop, in the menu bar, and on the phone at the gym.



22

WEEK CAMPAIGN, ONE
ADAPTIVE PLAN
MUTINY-LABS.COM

3

ADAPTIVE LAYERS, AI
TO MORNING GATE

100%

SELF-HOSTED, DATA
STAYS HOME

PRODUCTS

AI
COACHING

WEARABLES

// 01

The context

From the creators of ginoFit, the Garmin-connected health dashboard, Shanghai is the daily accountability layer for a 22-week run/walk campaign toward the Shanghai Marathon. The athlete trains by heart rate, not ego: every session is governed by a morning recovery verdict computed from overnight HRV, Body Battery and sleep. The app's one job is getting its athlete to the start line healthy.

// 02

The challenge

Nearly every fitness app rewards volume, so an athlete prone to overreaching gets cheered into injury. Shanghai had to do the opposite: reward restraint, cap what harder training can earn, and grade an easy day on a poor recovery morning as a win. The plan itself could not be static either, and all of it, private health telemetry included, had to live on the athlete's own hardware with no public exposure.

// 03

How we worked

We built a self-hosted system: a Python service on the athlete's own machine syncing wearable telemetry hourly, with a dark, installable web app served over a private encrypted mesh to phone and desktop. Generative AI grades each session rep by rep from sample-level heart data and can rewrite plan targets after every workout, inside hard safety rails, every change logged and revertible. And when the athlete disagrees, a conversational coach argues its case, or corrects bad sensor data on the record.

CLIENT

ginoFit · In-house product

SECTOR

Health & fitness

YEAR

2026

BUILT WITH

React progressive web app, Python API layer, on-device SQLite, wearable telemetry pipeline with sample-level heart rate analysis, generative AI coaching with realtime plan review and conversational data correction, generative art badge pipeline, Web Push, native macOS menu bar companion, private encrypted mesh networking

“Most training apps cheer you on while you dig yourself into a hole. Shanghai is the first one I have seen with the guts to tell an athlete to stop. That red gate is real coaching.”

EDWARD BETANCOURT, SHORT DISTANCE SPEED RUNNER AND MPT RECORD HOLDER

// 04

What we built

01 The morning gate

Overnight HRV, Body Battery and sleep combine into a green, yellow or red verdict each morning. It outranks every other number in the app, and obeying it is what the scoring system pays for.

02 AI-graded sessions

Every rep is rebuilt from sample-level heart data: pace, heart rate ceiling, cadence, and walk recovery judged on the true low reached, not the lap average. The AI coach then grades A+ to F in context: an easy session on a red morning is an A decision.

03 A coach that talks back

A chat coach with real authority. Ask why a session earned its grade and it cites the evidence. Tell it the treadmill over-read and it corrects the data with audited tools, re-grades on the truth, and tags the session as corrected.

04 A plan never set in stone

Three adaptive layers stack: the AI can rewrite block targets after each session inside clamped safety rails, a deterministic engine steps treadmill speed up or down from measured performance, and the gate has the final say. The result is one card with today's exact numbers.

05 Scoring that cannot be gamed

XP is grade times gate-correctness with a weekly cap, so junk miles can never beat disciplined ones. Twenty-three irezumi-style medals across seven categories, produced with a generative art pipeline, wait rusted and oxidized until earned: gold is minted only for the work actually done.

UNDER THE HOOD React progressive web app, Python API layer, on-device SQLite, wearable telemetry pipeline with sample-level heart rate analysis, generative AI coaching with realtime plan review and conversational data correction, generative art badge pipeline, Web Push, native macOS menu bar companion, private encrypted mesh networking

// 05

One training day, end to end

01

Overnight, the body reports in

Wearable telemetry syncs hourly to the athlete's own machine: HRV, Body Battery, sleep, and every session with sample-level heart rate. No cloud, no third party.

02

The gate rules the morning

At six the verdict lands as a push notification: green, yellow or red. The menu bar glyph tints to match, and if late data changes the color, the gate re-fires. Every other number defers to it.

03

One card of exact numbers

Block targets, the weekly schedule, the adaptive step and the gate combine into today's prescription. A 14-day projection shows what is coming and auto-tapers before race day.

04

The session is rebuilt rep by rep

Runs are scored on pace, heart rate ceiling, cadence and true walk-recovery lows. Strength work is parsed set by set with its own rubric. Sensor glitches are excluded and cardiac drift detected.

05

The AI passes judgment

A grade from A+ to F, argued in context of the morning's gate. Then the plan review decides whether the block's targets should move, inside clamped rails, every change logged and revertible.

06

The athlete gets the last word

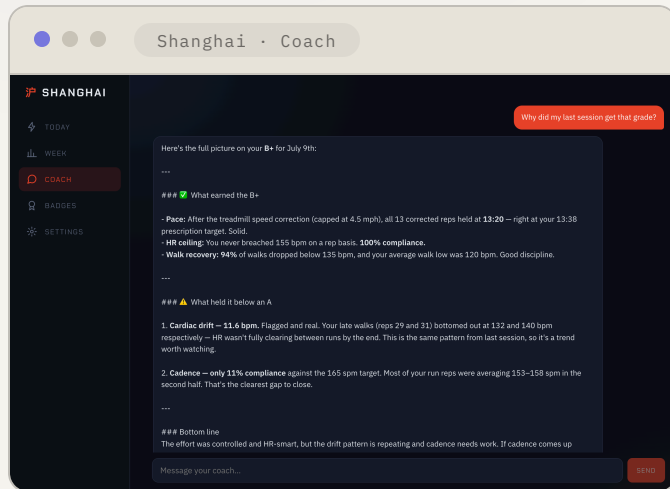
The chat coach explains any verdict, corrects bad sensor data on the record, and re-grades on the truth. XP, medals and notifications close the loop for the day.

UNDER
THE
HOOD

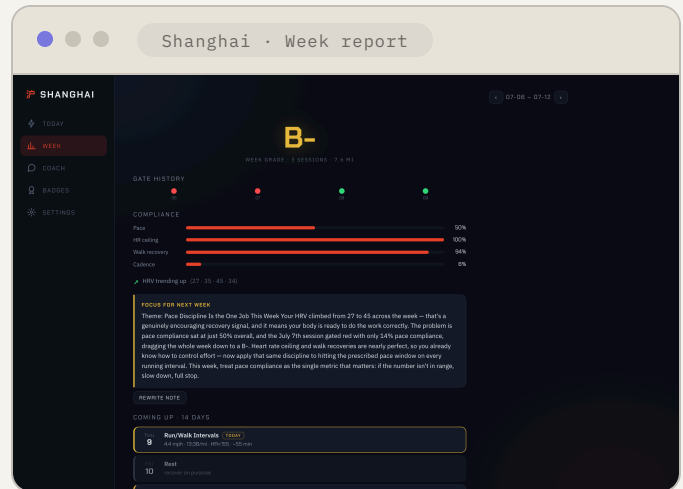
React progressive web app, Python API layer, on-device SQLite, wearable telemetry pipeline with sample-level heart rate analysis, generative AI coaching with realtime plan review and conversational data correction, generative art badge pipeline, Web Push, native macOS menu bar companion, private encrypted mesh networking

// 06

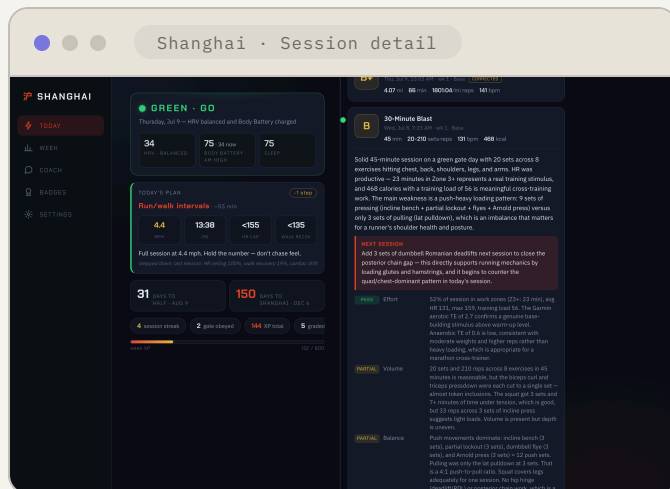
Inside the product




Asked to defend a grade, the coach cites the corrected treadmill data, a perfect heart rate ceiling and strong walk recoveries, then names cadence as the gap to close.



The Sunday report card: a week grade, gate history, compliance bars per rule, and an AI focus note that assigns the week ahead exactly one job.



Strength is a first-class citizen: sets, reps and heart rate zones graded on their own rubric, a push-heavy imbalance called out, and the fix prescribed for the next session.



Twenty-three irezumi-style medals across seven categories, from First Blood to the finisher's dragon. Locked medals sit rusted; gold is minted only when earned.

// 07

The campaign, by the numbers

22

WEEK CAMPAIGN, ONE ADAPTIVE PLAN

3

ADAPTIVE LAYERS, AI REVIEW TO MORNING GATE

23

MEDALS IN SEVEN CATEGORIES, GOLD ONLY WHEN EARNED

4


SURFACES: PHONE, DESKTOP, MENU BAR, NOTIFICATIONS

100%

SELF-HOSTED, TELEMETRY NEVER LEAVES HOME

0

PUBLIC PORTS, PRIVATE ENCRYPTED MESH ONLY



ON THE PHONE AT THE GYM: A RED MORNING'S ORDERS, AND THE MEDAL WALL WITH DISCIPLINA EARNED FOR OBEYING THEM

// 08

What this demonstrates

○ AI-native product design

○ Wearable data engineering

○ Self-hosted private systems



• SHANGHAI • RACE DAY, DECEMBER 2026

22

WEEK CAMPAIGN, ONE ADAPTIVE PLAN

3

ADAPTIVE LAYERS, AI TO MORNING GATE

100%

SELF-HOSTED, DATA STAYS HOME

The hardest thing to coach is restraint. Shanghai makes it the highest-scoring move in the game.

ABOUT MUTINY LABS

Mutiny Labs is an innovation and strategy studio in Puerto Rico. We build digital infrastructure for organizations that move the island forward, designed with the people who use it and built to be run by the teams that own it.

Products & Platforms

Web applications, portals and operations software shaped around real workflows.

Applied AI

Generative and analytical AI placed deliberately, private by construction.

3D & Immersive

Virtual museums, gaussian-splat capture and experiences that carry culture.

Data & Measurement

Impact platforms, dashboards and reporting that funders can trust.

Let's build yours. mutiny-labs.com