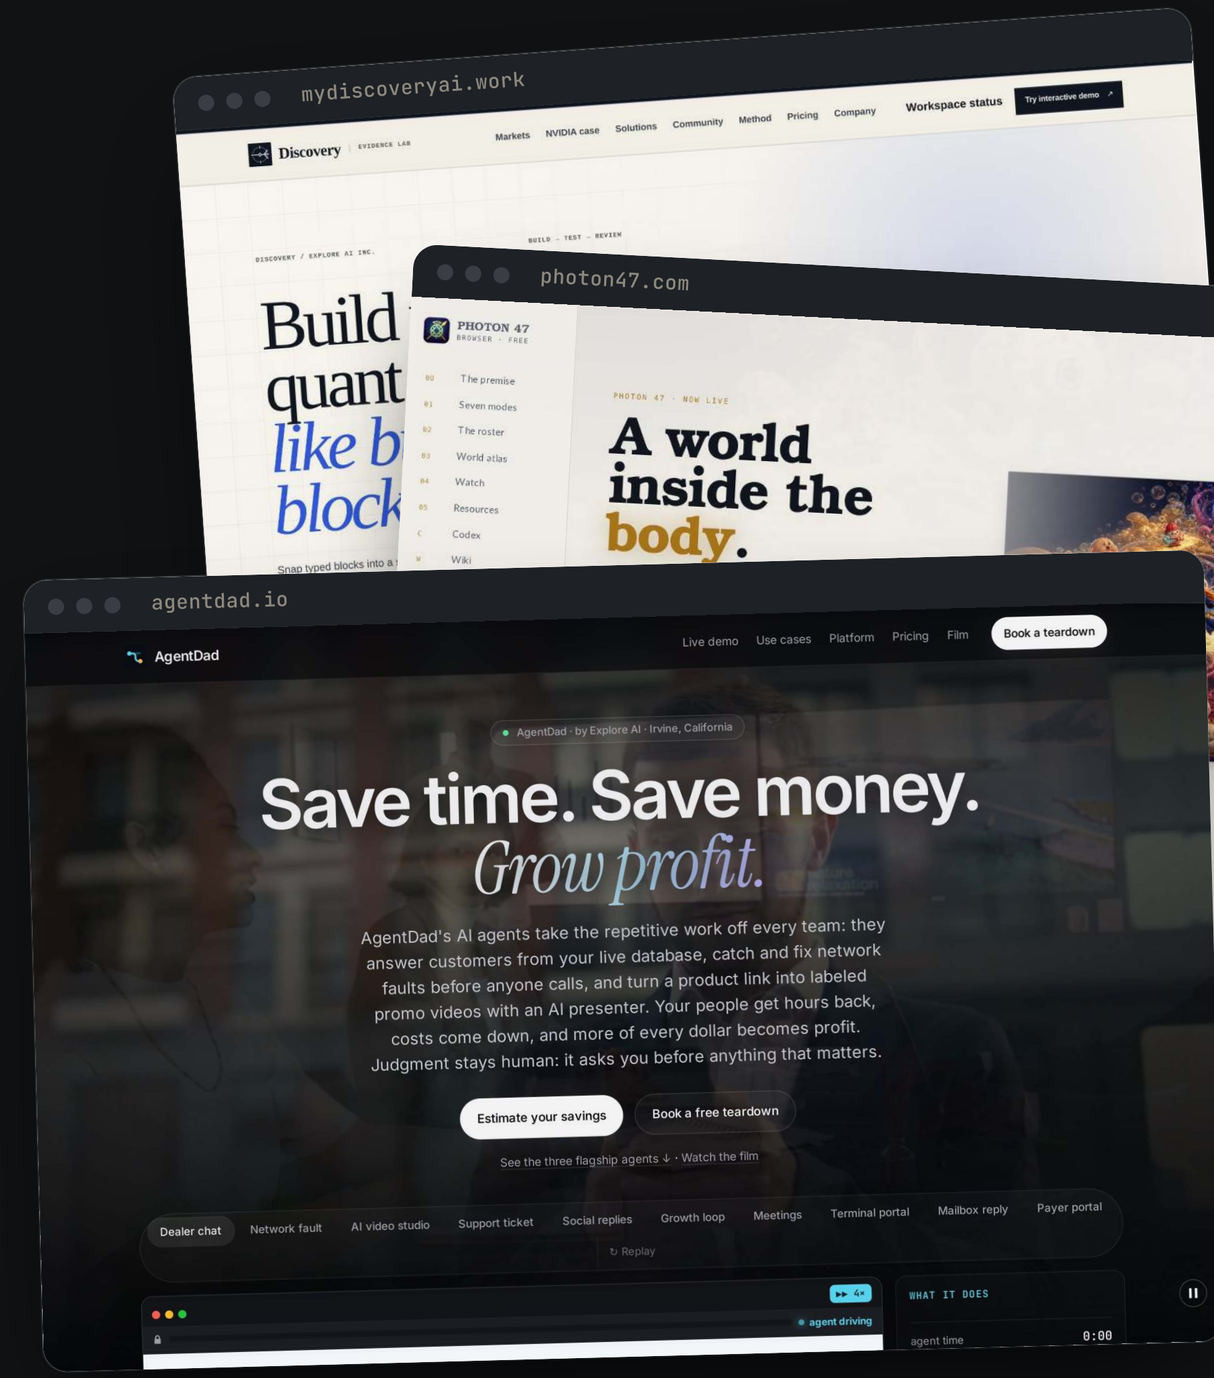
explore.AI

COMPANY OVERVIEW · 2026

Agents that do the work. *Models you own.*

Explore AI builds supervised AI agents that save companies time and money and grow profit: running operations today, building and shipping software next, on a coding model the company controls.



EXPLORE AI INC. IRVINE, CALIFORNIA EXPLORE-AI-INC.COM

THE COMPANY IN ONE SLIDE

We turn AI into *profit*: agents that save time and money in a company's operations today, ship its apps next, and run on a model it owns.

30 *days*

Teardown to a production agent

AgentDad's published promise, with pricing from a $15,000 pilot to a $200,000 Premium deployment.

1 → *5*

One codebase, five storefronts

Our own game ships from one source to web, Steam, CrazyGames, Poki and itch.io: the core of App Factory.

≤32*B*

Sovereign coding model

The best agentic coding model that fits one server, deployed and improved inside the customer's walls.

THE PROBLEM

AI can write. It still can't be trusted *to finish the job.*

01

Operations still pay people to retype.

Invoices checked against rate confirmations, referrals keyed in from faxes, purchase orders retyped into the ERP. Each step is simple; the volume and the switching cost whole teams their week, and the payroll that goes with it.

02

Shipping software still takes a team.

A real product needs a frontend, a backend, login, a database, payments, tests and store review. Code generators stop at the prototype; the last mile is where months go.

03

Companies can't rent their brain.

Regulated firms and governments need models that run inside their own walls. Sovereignty requirements could shape **30–40% of AI spending** by 2030.[1]

[1] McKinsey estimate, as reported; see Sources.

OUR SOLUTION

One agentic engine. *Three compounding businesses.*

01 · RUN OPERATIONS

AgentDad

Save time and money, grow profit

Supervised agents inside the systems a company already uses: they answer from its records, fix what policy allows, and ask a person before anything risky.

LIVE · SELLING

02 · BUILD & SHIP APPS

App Factory

Describe an app; agents ship it

Agents build the full stack (frontend, backend, login, database, payments) and publish with one click to the App Store, Google Play, Steam and the web.

NEXT · IN DEVELOPMENT

03 · OWN THE MODEL

Sovereign Coder

≤32B agentic coding model

The best coding model that fits one server, deployed on-premises or in a private cloud, then continuously improved on the customer's own tasks.

BUILDING · IN TRAINING

SHARED ENGINE **Read → Search → Act → Ask**

Approval gates · traced, replayable runs
verifiable rewards · distillation into small models

PROOF WE SHIP, TODAY Photon47 & studio games · Discovery financial analysis · OnlineNIW

WHY NOW

Three shifts make this *the moment.*

Agents can now act reliably inside real systems, small models are good enough to own, and buyers increasingly demand control over both.

SHIFT 01

Agents moved from demo to budget line.

46%

Annual growth of the AI agents market: $7.8B in 2025 to $52.6B by 2030.[2]

SHIFT 02

Building software became a consumer product.

$400*M*

ARR at Lovable, an AI app builder, in February 2026.[3] Generating code is solved. Shipping and operating it is not.

SHIFT 03

Sovereignty became a requirement.

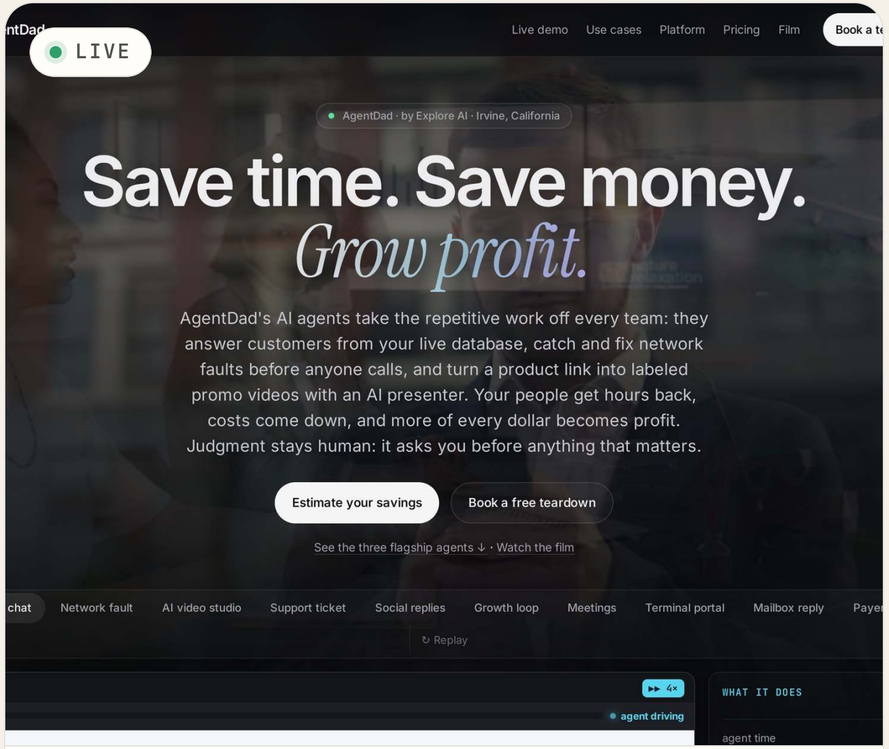
30–40%

Share of AI spending that sovereignty requirements could shape by 2030.[1] Small, strong models make it deployable.

[1] McKinsey, as reported · [2] MarketsandMarkets · [3] TechCrunch. Full citations in Sources.

LIVE TODAY

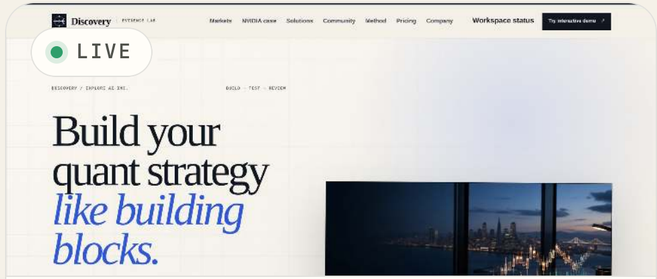
Four products and a game studio, *already in market.*



AgentDad

B2B operations agents for support, sales, marketing and back office.

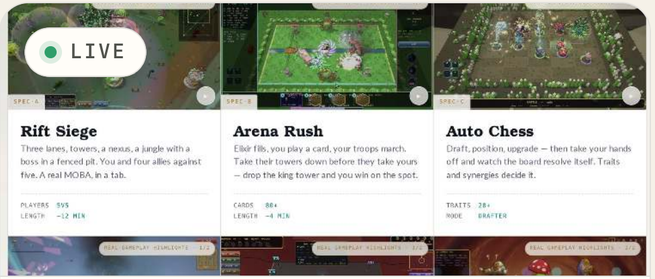
`agentdad.io`



Discovery

Quant strategies built like blocks.

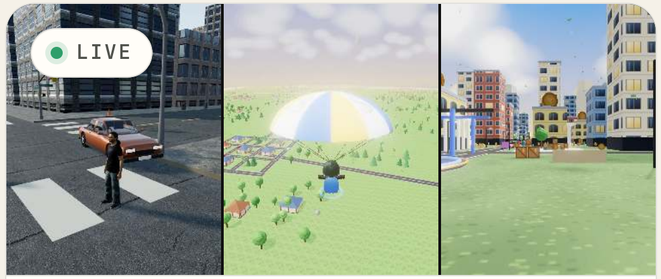
`mydiscoveryai.work`



Photon47

Multiplayer 3D game, seven modes.

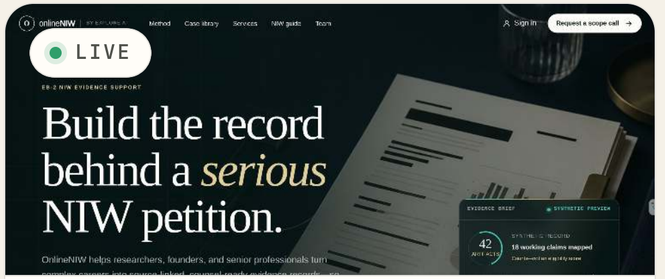
`photon47.com`



Studio games

LIYF open city · Flag Fight royale.

`explcre.github.io`



OnlineNIW

EB-2 NIW evidence service.

`online-niw.com`

PRODUCT 01 · AGENTDAD · B2B

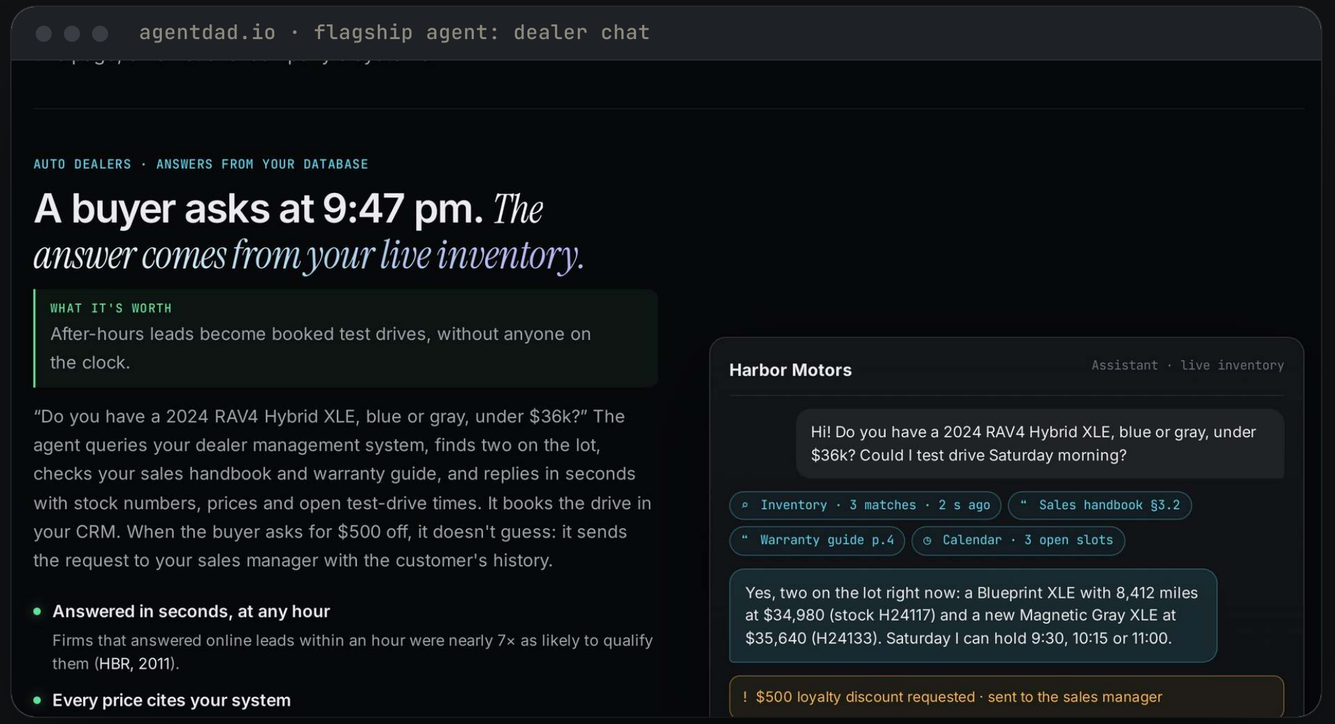
Save time. Save money. *Grow profit.*

1 **Read**
Pull the fields that matter from invoices, faxes, contracts, chats and alerts.

2 **Search**
Answer from live databases, policies and past tickets, citing the source every time.

3 **Act**
Work existing systems through APIs, or supervised browser actions where none exist.

4 **Ask**
Judgment stays human: it stops for a person before money moves or a customer commitment.



DAY 0	DAYS 1–10	DAYS 11–17	BY DAY 30
Free 30-minute workflow teardown	Connectors, rules, golden test set	Shadow mode, scored against the team	Autopilot with per-action approvals

Every agent is judged by one number: *your profit.*

AgentDad starts where an agent clearly pays for itself, and shows the math before the customer spends a dollar. Priced against the work it replaces, not per seat.

TIME BACK

2,851 *h*

Hours a year

The team stops keying, checking and chasing, and spends the time on customers and growth.

MONEY SAVED

$74,304

Net in year one

After AgentDad's fees, for a single workflow.

PAYBACK

4 *months*

To pay for itself

Then the savings compound as more workflows go live.

REVENUE UP

7×

Likelier to qualify a lead

When it is answered within the hour, as agents do at any hour.[5]

5,616 h of manual work a year → **2,851 h** handled by one agent → **$128,295** of labor at $45/h → **$74,304** net after our fees

Worked example from AgentDad's published savings calculator: six people spending 18 hours a week each on one workflow, at $45 an hour, with one agent. Not a measured client result; each teardown measures the customer's own numbers. [5] Harvard Business Review, 2011.

Land one workflow. *Expand across departments.*

$0	**Teardown** One workflow mapped; fixed written quote in 2 business days.
$15,000	**Pilot** Fixed scope, one workflow live within 30 days; credited to Production.
$4,500 /agent/mo	**Production** Annual; hosting, models, 24/7 monitoring, monthly evals.
$200,000	**Premium · 90 days** Up to five workflows across departments with a dedicated pod.
Custom	**Enterprise** Private cloud or on-premises, SSO, SLAs: the bridge to Sovereign Coder.

WHERE WE START

Logistics · Healthcare · Insurance · Property management · Lending · Manufacturing · Law firms · Auto dealers

126
Researched target companies in Southern California, Irvine first.

54
Mapped use cases; 8 interactive demo cases on the site.

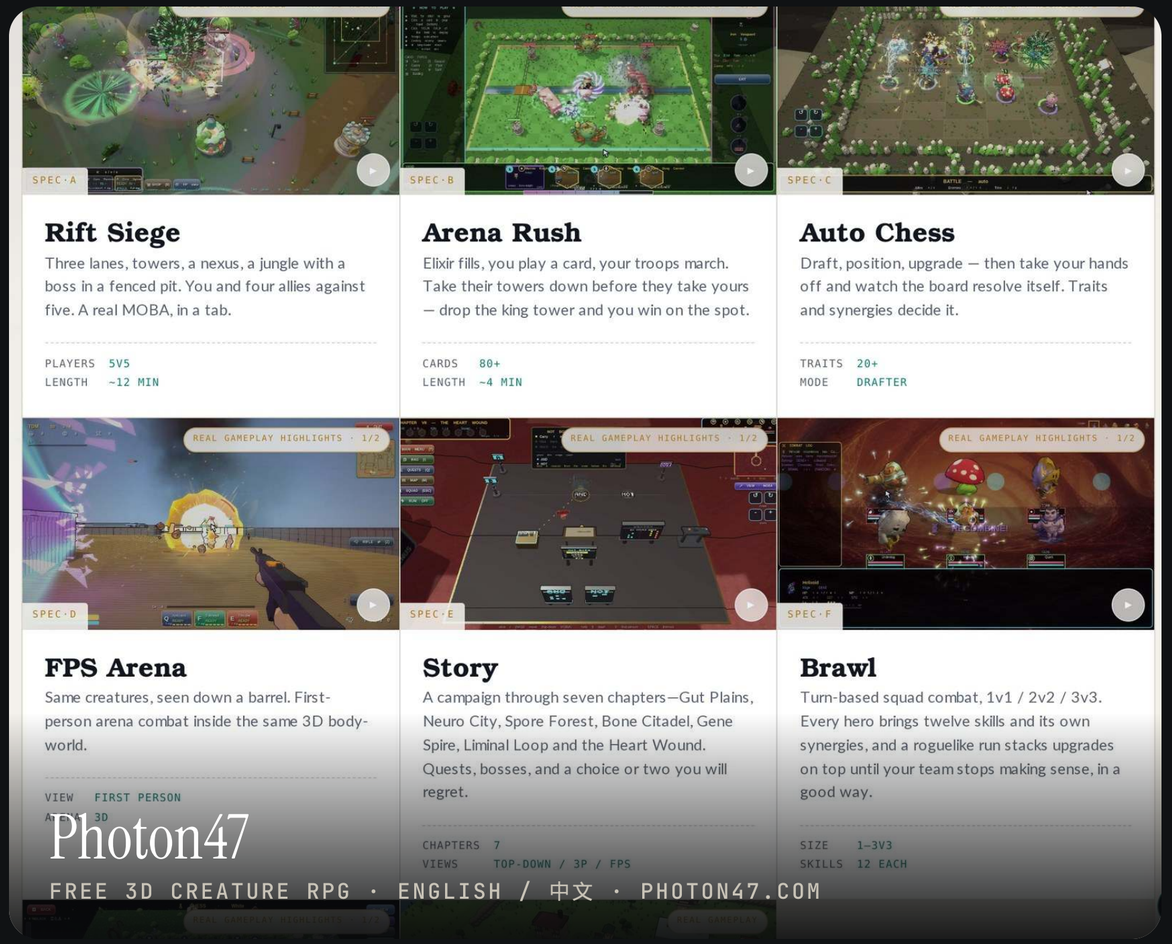
2
Product films, English and Mandarin, showing nine workflows.

0
Training on customer documents. Deploy in our cloud, theirs, or on-premises.

Pricing as published on agentdad.io (October 2026). Target-company list is an outreach pipeline, not customers.

PRODUCT 02 · GAME DEVELOPMENT

A game studio that ships *at AI speed.*



7

Game modes on one roster: MOBA, deck battler, FPS, auto chess, story, brawl, chess.

45

Playable heroes plus 6 bosses, with cross-play PvP and cloud saves.

5

Storefronts from one codebase: web, Steam, CrazyGames, Poki, itch.io.

2,000+

Commits behind ~600 automated release checks, Stripe payments, Postgres backend.





> Our release pipeline is *App Factory's prototype.*

Build your strategy like blocks. *Let an AI agent trade it.*

Snap strategies together LIVE

25 typed blocks wire into a strategy graph; every edit recomputes the chart, fills and backtest instantly.

Your AI trading agent COMING NEXT

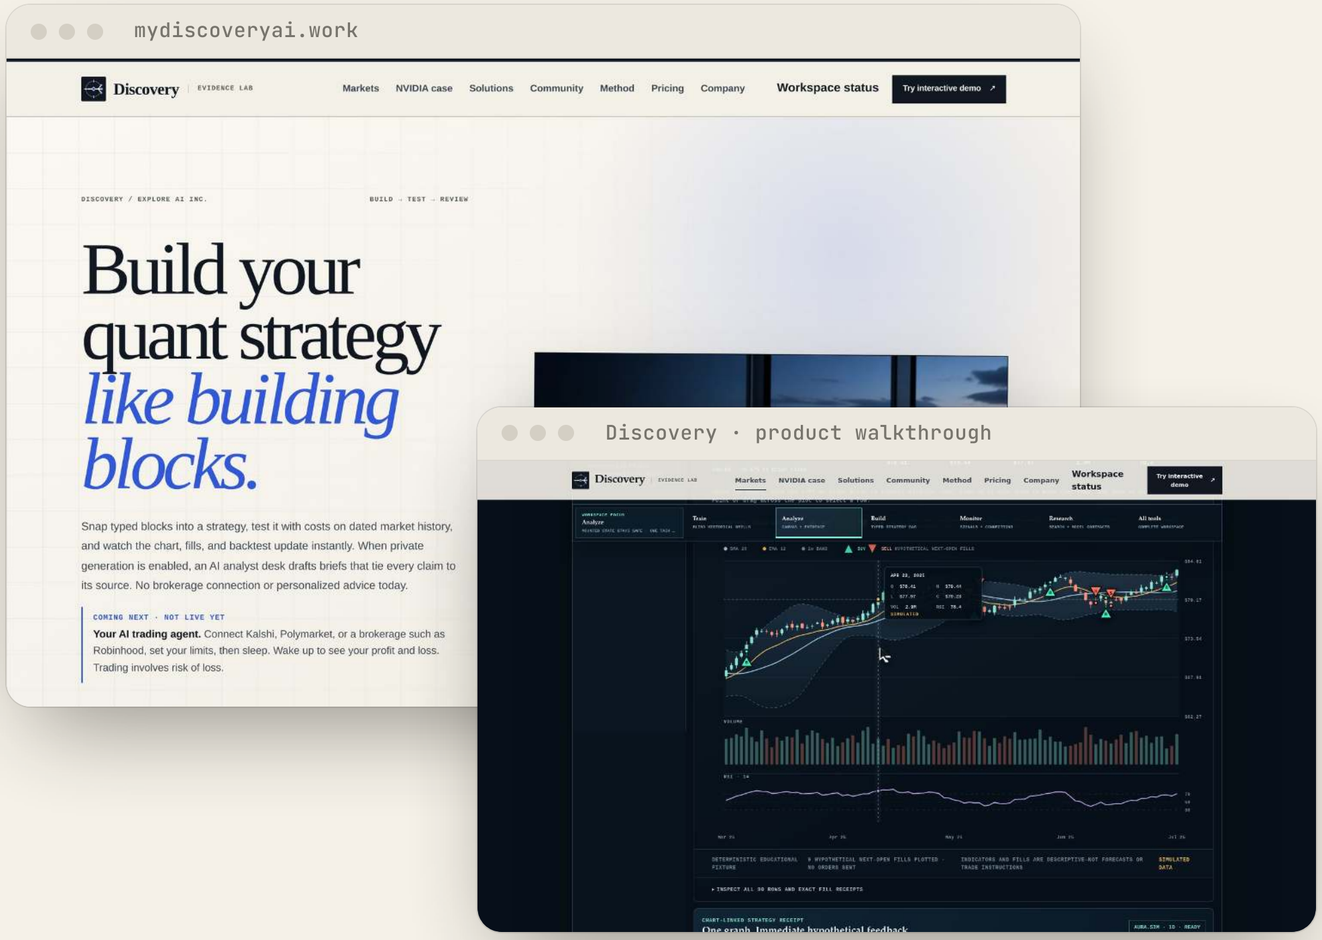
Connects to Kalshi, Polymarket and brokers like Robinhood via MCP. Set your limits, then sleep; **wake up to see your profit and loss.**

An AI analyst desk PRIVATE TESTING

Agents draft research briefs that keep observations, scenarios, invalidation points and sources separate.

Test before you trust LIVE

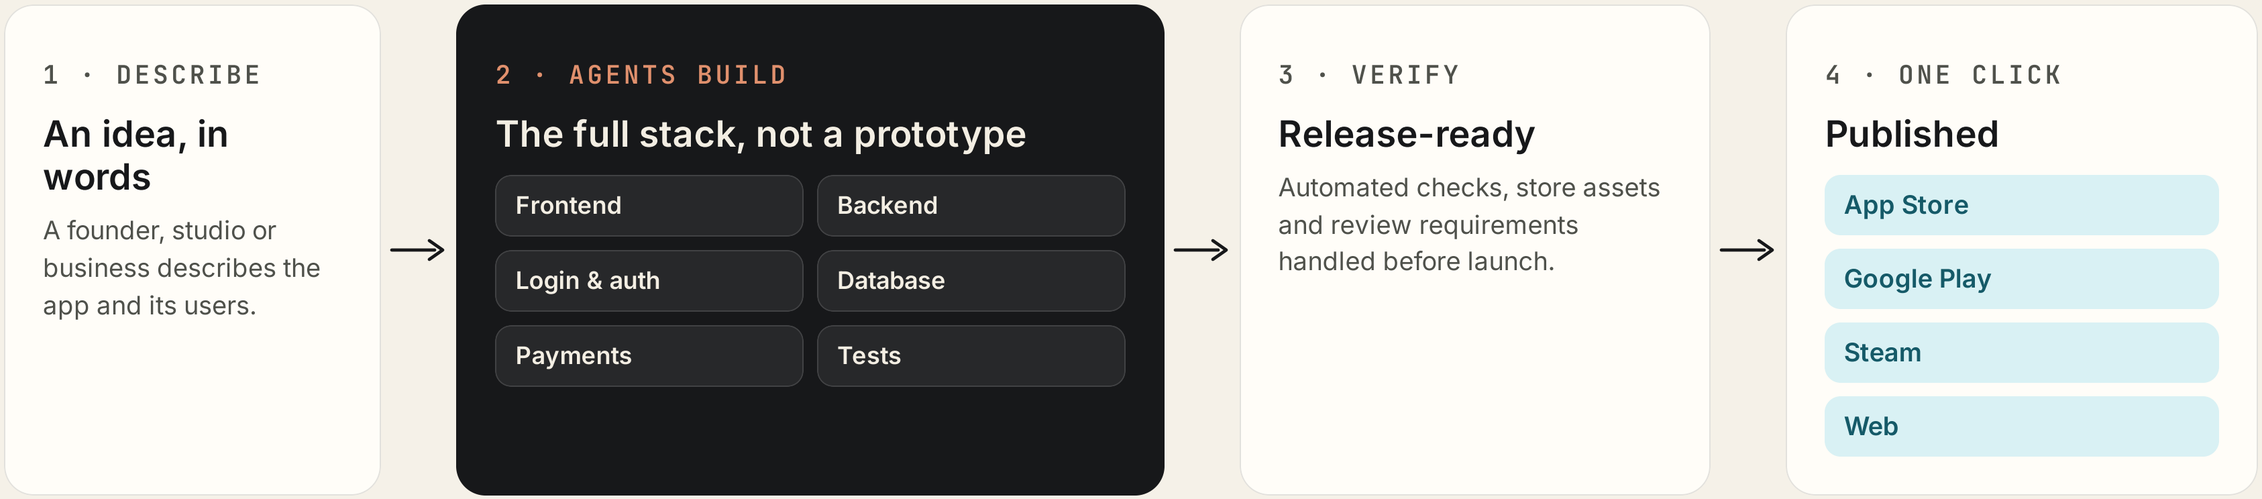
Bounded search with holdout gates, and drills that reveal the next candle only after you decide.



NEXT · APP FACTORY

Describe it. *Ship it everywhere.*

Our aim: an idea becomes a shipped, revenue-ready app in days instead of months, without hiring a team to build it.



We already run this pipeline.

Photon47 ships from one source to five storefronts behind ~600 automated checks, with auth, Postgres and Stripe.

Agents that operate, not just generate.

AgentDad's approval gates and traced runs carry over to deployment, store review and live operations.

Owned-model economics.

Built on our Sovereign Coder, so every app gets cheaper to build as the model improves.

BUILDING · SOVEREIGN CODER ≤32B

The best agentic coding model *that fits in one server.*

Runs on a single GPU node

Lower, predictable cost: on-premises or private cloud, no per-token API bill.

Code and data never leave

Built for regulated industries, critical infrastructure and governments.

Improves on the customer's own work

Adapted to their codebase and tasks with verifiable rewards and distillation.

Deploy → Adapt → Improve ↺

RESEARCH EDGE · FOUNDER'S WORK

57.0%

SWE-bench Verified, 429 held-out tasks. Consensus-plus-judge selection over three attempts vs 48.9% for one attempt; pre-registered, $p \leq 10^{-4}$.

+15.8%

WebShop, 1.5B student. UOPD distillation scores 76.8 vs 66.3 for on-policy distillation from the same 7B teacher.

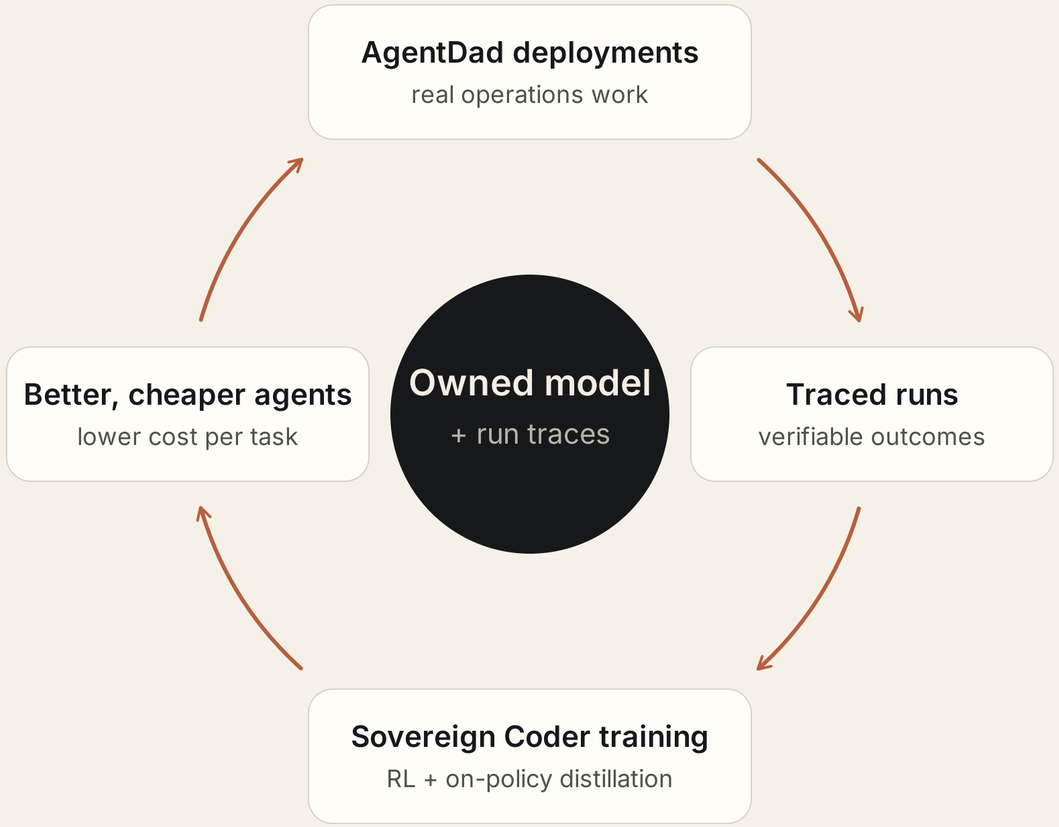
28,750×

Smaller distilled experts that keep 95% of teacher performance on average across five teachers.

Research results from the co-founder's manuscripts (arXiv:2609.34036 and papers under review), not product benchmarks.

THE FLYWHEEL

Every deployment sharpens the model. *Every better model widens the moat.*



Deployments create signal. Each AgentDad workflow and App Factory build produces a traced run with a checkable outcome.

Signal trains the model. Verifiable rewards and distillation turn those traces into a smaller, stronger Sovereign Coder.

A better model wins the next deal at lower cost per task, and becomes a product on its own.

Customer documents are never used to train shared models. The flywheel learns from our own runs, synthetic tasks and opt-in programs.

BUSINESS MODEL

Services revenue today. *Software and model margins next.*

PRODUCT	CUSTOMER	HOW WE EARN	STATUS
AgentDad	Operations, support, sales and marketing teams	$15k pilots → $4.5k/agent/month production → $200k Premium; enterprise contracts	LIVE
App Factory	Founders, SMBs, agencies, game studios	Subscription plus usage; managed store publishing	BUILDING
Sovereign Coder	Enterprises, regulated industries, public sector	Model license, deployment fee, improvement retainer	TRAINING
Games	Players worldwide, English and Chinese	Free-to-play: in-game purchases, ads, storefront releases	LIVE
Discovery	Individual investors and research teams	Subscription tiers; AI trading agent plans (coming next)	LAB LIVE
OnlineNIW	Researchers, founders, senior professionals	Fixed-scope evidence services, from $495	LIVE

MARKET

Large markets. *One engine across all of them.*

Each product line sits in a market measured in tens to hundreds of billions of dollars, and all of them run on the same agents and models.

$500–600*B*

Sovereign AI by 2030

Infrastructure, models and applications shaped by sovereignty requirements.[1]

→ SOVEREIGN CODER

$234*B*

Global games by 2028

Up from $201.6B in 2025, the first year above $200B.[4]

→ STUDIO · APP FACTORY

$52.6*B*

AI agents by 2030

From $7.8B in 2025, growing 46% a year.[2]

→ AGENTDAD

$400*M*

One app builder's ARR

Lovable in February 2026: proof that people pay to turn ideas into software.[3]

→ APP FACTORY

[1] McKinsey, as reported · [2] MarketsandMarkets · [3] TechCrunch · [4] Newzoo. Independent estimates; see Sources.

COMPETITION

Others generate, automate or rent.
We finish and let you own it.

	Chat copilots	RPA suites	AI app builders	Frontier model APIs	Explore AI
Acts inside existing systems, with approval gates	◐	●	○	○	●
Answers from company records, with citations	●	○	○	◐	●
Builds full-stack apps: backend, login, database	◐	○	●	○	◐
Publishes to App Store, Google Play and Steam	○	○	◐	○	◐
A model the customer owns, on-premises	○	○	○	○	◐
Improves on the customer's own tasks	◐	○	○	◐	●

● offered · ◐ partial, or on our roadmap for Explore AI (App Factory and Sovereign Coder are in development) · ○ not offered. Category-level view; individual vendors vary.

TRACTION & PROOF

Small team. *Unusual output.*

PRODUCTS

5

Live product sites

agentdad.io, mydiscoveryai.work, online-niw.com, photon47.com, explore-ai-inc.com.

AGENTDAD

30*d*

Priced, published offer

Teardown to production in 30 days; $15k pilot to $200k Premium.

PIPELINE

126

Target companies researched

Southern California, Irvine first, with tailored outreach drafted.

USE CASES

54

Mapped workflows

8 interactive demo cases and two product films (English, Mandarin).

GAMES

7 *modes*

45 heroes · 5 storefronts

Photon47, built solo: multiplayer, payments, cloud saves.

ENGINEERING

2,000+

Commits on Photon47

With ~600 automated checks gating every release.

RESEARCH

ICML · KDD

2026 papers (co-author)

Plus ACM BCB 2024 SIGBio Best Paper (co-author) and UOPD (co-first).

DISCOVERY

3:32

QA-verified walkthrough

Public market lab, evidence case and trainer, live without an account.

Figures from our live sites and repositories (October 2026). The company is early; we report no revenue figures here.

TEAM

A scientist who ships, *and an operator who builds companies.*

Pengcheng (Ryan) Xu

CO-FOUNDER · RESEARCH & TECHNOLOGY

UC Irvine CS PHD · **AWS** SOFTWARE ENGINEER · **Intel** DL SOFTWARE ENGINEER · **XtalPi** RESEARCH SCIENTIST · **Pandita AI** LLM RESEARCH SCIENTIST

EDUCATION

PhD Computer Science, UC Irvine (GPA 4.0); MS UIUC (AI@UIUC); BS SJTU–Michigan Joint Institute, Outstanding Graduate.

RESEARCH

Co-first author of UOPD (agent distillation); co-author at ICML 2026, KDD 2026; ACM BCB 2024 SIGBio Best Paper.

ENGINEERING

AWS (2023 internship): VMware Cloud on AWS console and Lambda backend. Intel (2021–22 internship): Neural Compressor quantization, C++ inference server.

BUILDER

Built Photon47 solo: seven modes, multiplayer, payments, five storefronts. Built the AgentDad and Discovery products.

CO-FOUNDER

Company building & operations

More than 20 years of building and operating businesses. Leads company development, operations and long-term planning.

HIRING NEXT

- Founding ML engineer · RL & distillation
- Full-stack & mobile engineer
- Forward-deployed agent engineer
- Go-to-market lead

PLAN · NEXT 18 MONTHS

Three milestones, *each a step up in value.*

MONTHS 0-6

Prove AgentDad

- Paid pilots converted to production workflows
- App Factory alpha ships our own apps and games
- Sovereign Coder v0 on internal agentic coding evals

EXIT CRITERION

Pilots renewing into paid production workflows.

MONTHS 6-12

Launch App Factory

- Public beta: full-stack app with one-click store publishing
- Multi-department Premium deployments
- Sovereign Coder v1 (≤32B) with public benchmark results

EXIT CRITERION

Customer apps live in stores, built end to end by agents.

MONTHS 12-18

Sovereign deployments

- First on-premises deployments with improvement retainers
- App Factory paid tiers and studio partners
- Flywheel live: traced runs feed model updates

EXIT CRITERION

A sovereign deployment under contract, improving on the customer's tasks.

Targets, not commitments; timing depends on funding and hiring.

explore.AI

Agents that do the work.
Models you own.

SAVE TIME · SAVE MONEY · GROW PROFIT

EMAIL
contact@explore-ai-inc.com

COMPANY
explore-ai-inc.com

PRODUCT
agentdad.io

APPENDIX A · RESEARCH FOUNDATION

The science behind *small, capable agents.*

AGENT DISTILLATION · UOPD

Uncertainty-aware on-policy distillation

The teacher steps in only where it is uncertain. Best ALFWorld success at 3B and 1.5B students; WebShop 76.8 vs 66.3; 13.8–36.5% less training time than a strong baseline.

+15.8%

CODING AGENTS · RL

Spending the whole rollout budget

Routes each GRPO group by its failure mode for repository-level coding. Selection lifts SWE-bench Verified resolution to 57.0% from 48.9% (429 held-out tasks, pre-registered).

+8.1 *pts*

DISTILLATION AT SCALE · OMEGAGENOME

Sub-million-parameter experts

Two-stage distillation from large genomic language models: experts up to 28,750× smaller retain 95% of teacher performance on average.

28,750×

ICML 2026

A Very Big Video Reasoning Suite

Co-author · video reasoning data and generators

KDD 2026

MUGO: causal variant discovery

Co-author · differentiable combinatorial optimization

ACM BCB 2024

L2 normalization for single-cell visualization

Co-author · SIGBio Best Paper Award

ACL BIONLP 2024 · 2025

LLMs for biomedical lay summaries

Co-author, then lead author

UOPD: arXiv:2609.34036 (co-first author; under review). Coding-agent and OmegaGenome results are from manuscripts under review. Research was carried out at UC Irvine with collaborators.

Sources

[1] **Sovereign AI.** McKinsey & Company estimate that sovereignty requirements could shape 30–40% of AI spending, a $500–600B opportunity by 2030, as reported in secondary coverage (2026).

[2] **AI agents market.** MarketsandMarkets, "AI Agents Market": $7.84B (2025) to $52.62B (2030), 46.3% CAGR.

[3] **Lovable ARR.** TechCrunch: Lovable crossed $400M in annual recurring revenue in February 2026, confirmed by the company.

[4] **Games market.** Newzoo, Global Games Market Report: $201.6B in 2025; $234.4B projected for 2028.

[5] **Lead response.** Harvard Business Review (2011), "The Short Life of Online Sales Leads": firms that contacted leads within an hour were nearly 7× as likely to qualify them.

AgentDad savings example. agentdad.io calculator: six people, 18 hours a week each on one workflow, $45 an hour, one agent.

Product facts. agentdad.io pricing and deployment terms; mydiscoveryai.work; online-niw.com; photon47.com; explore-ai-inc.com (all as of October 2026).

Research. arXiv:2609.34036 and the co-founder's manuscripts under review; see Appendix A.

This presentation is for discussion only and is not an offer to sell or a solicitation of an offer to buy any security. It contains forward-looking statements about plans and products in development (App Factory, Sovereign Coder) that involve risks and may change. AgentDad demonstrations use fictional companies and simulated runs. OnlineNIW is not a law firm and does not provide legal advice. Discovery provides no personalized investment advice and does not execute trades today; agentic trading is a roadmap item, and trading involves risk of loss. Third-party names are used for identification only and do not imply endorsement.