

SEAN ACRES

seanacres@gmail.com · 315.323.0854 · San Francisco, CA · linkedin.com/in/seanacres

SUMMARY

- **AI product leader who builds: agent harnesses, generative pipelines, and the evals that decide whether the behavior is good enough**
- Founding PM to VP of Product: seven 0-to-1 launches, a #1 iOS game, a title on a 180M-subscriber franchise, a 3.6M-MAU org
- Ships at engineering-org scale by directing fleets of coding agents, then turns the method into shared rails the whole team runs on
- Owned P&L at three companies; hired and led product, design, and growth orgs across the US and India, from seed through Series C
- Voice AI since 2017; ships production Python, TypeScript and Swift and the evals behind them; spec and prototype from one person

EXPERIENCE

Weekend (formerly Volley), AI Voice and Entertainment, \$55M Series C | *Senior Product Manager* Jul 2024 - Present

Foundry, an AI agent harness that builds shippable games

- Owns the product and technical roadmap and the quarterly stakeholder commitments five cross-functional teams plan against
- Built a multi-tenant game runtime in nine days; the platform took it as the starting point over migrating the existing server code
- Shipped the generative art pipeline that replaces placeholder art across Foundry; six teammates took access in week one
- Replaced a deterministic 45-question flow with one adaptive AI interview, deleting the legacy paths instead of flagging them
- Stood up the service that generates game content end to end: an independent LLM judge screens it, replacing self-reported safety
- Wrote the shared build system every Foundry repo runs on, and cut the platform's total CI time by 60%
- Improved the harness version after version; the studio's output went from months per title to five launches in five weeks
- Produced hundreds of games through the harness so far, at a rate still climbing; 10+ finished titles sit in the launch queue

A game-show title in development

- Built an AI game show host end to end in two weeks; 40+ playtesters named it the standout feature of the game
- Wrote evals covering every dialogue scenario, with an LLM judge grading the runs and self-tuning the prompts behind them
- Automated the question bank with agents that generate and grade content: 1,056 questions regraded, difficulty corrected on 2,541
- Prototyped a parallel scaffold for generating any game on the production stack, ahead of the harness that productionized first

Company-wide AI and engineering

- Leads AI adoption: the onboarding path, the shared agent rails, and internal plugins for orchestrating agents, review, and research
- Created the vibe-coding curriculum 80+ employees took before a hackathon; a team with no engineers shipped a voice-controlled game

Games and platform (2024-25)

- Turned a planned 2D overlay into a 3D game on Moonbug's flagship IP (180M+ subs); five company firsts in 10 months, now the template
- Recruited the external 3D art team that gave the company a capability it did not have, and led five people internally
- Wrote the case for cloud streaming and aligned every stakeholder; CoComelon went from one platform to three and hit the year-end goal
- Built a COPPA-compliant kids' voice prototype to settle a roadmap question without spending any engineering time
- Owned Wheel of Fortune on Roku, one of two TV titles: 10 experiments in two quarters, six positive, an estimated \$384K in ARR
- Ran a team hackathon that turned members' own game ideas into three production experiments in two days

Joyride Games (mobile games studio) | *Staff Product Manager (Contract)* Aug 2023 - Jan 2024

- Grew an animated character IP from zero to 150K followers and 3M views in five months across TikTok and YouTube Shorts
- Set the weekly experiment cadence behind it, treating every post as a test with a read-out rather than a publication

Fractal Capital | *Staff Product Manager & Venture Partner* May 2021 - Aug 2023

- Founded both companies out of the firm as its venture partner, taking them from incorporation to launched product
- Built Convo.cash, a conversational bot builder shipped pre-ChatGPT, and Hype.day in web3: 60K users, partners including Coinbase
- Pitched investors and raised a \$350K seed for the web3 product, and owned the P&L across both companies
- Led product strategy across the web3 and creator-economy verticals, the two markets the firm was building into

Rabbit (a16z-backed, acquired) | *VP of Product* Oct 2018 - May 2019

- Ran the product org for a Series B streaming startup at 3.6M MAU; hired and led 3 PMs and designers across 20+ engineers
- Led a full web and mobile redesign, reorienting the product around permanent groups and servers instead of one-off rooms

Joyride (voice AI) | *Director of Product* May 2017 - Jun 2018

- Founding PM on a Google speech-to-text voice prototype in 2017; pivoted to live trivia and hit #1 in Trivia on iOS with 1M+ downloads
- Hired and led the US product, design and growth teams, standing up design and growth as functions the company did not have
- Directed product and engineering out of the India studio through launch, managing its PMs and the studio's operations

Whisper (now MediaLab) | Lead Product Manager

Apr 2016 - Apr 2017

- Grew web MAU 300%+ in 10 months into the top 200 mobile sites; improved ad CTR 50% and cost per impression 20%
- Built an ML article generator with the data-science team in 2016, turning user posts into content for the L.A. Times
- Led product on the ads infrastructure behind that growth; it is still the core revenue system the company runs on today

Wave | Founder

May 2015 - Apr 2016

- Co-founded an ephemeral iOS messaging app; learned Objective-C to build features instead of queueing behind an engineer

GREE International | Senior Product Manager

Sep 2014 - May 2015

- Owned P&L on a top-25 grossing title: rebalanced the game economy and merchandising, lifting LTV 35% and \$300K+ monthly revenue
- Managed a 20+ person mobile game studio across product, design, engineering, QA and game design

Kiwi, Inc. (acquired by Jam City) | Director of Product

Apr 2013 - Sep 2014

- Designed the first casual mobile game to ship guilds and seasons, an approach presented at GDC Next 2014
- Shipped the approach across multiple top-100 grossing Android titles; Homescapes and Gardenscapes adopted it years later
- Owned product for 100% of company revenue across top-25 grossing titles; led a 40-person studio inside a 160-person company

PROJECTS

Rareta (rareta.com) | Autonomous car-buying agent

Name a car and a ceiling; it hunts dealer inventory around the clock and negotiates by email while you do nothing

- Pulled live inventory straight from manufacturer and used-car APIs, so the ranking saw the real market, not one dealer group's lot
- Wired every source behind one adapter interface, so covering a new brand is a config entry rather than a rewrite
- Scored every listing on what it would actually cost delivered, not the sticker, and emailed the dealers holding the best
- Tested it against every kind of reply a dealer can send: nothing wrong went out, and nothing right was held back
- Blocked its own email whenever the model invented a price or drifted off-script, sending nothing rather than something generic

Tellspike (tellspike.com) | On-device AI assistant for iPhone

Ask it anything without the question leaving your phone; when it does reach the cloud, it sends an anonymized rewrite

- Built local memory on the animachord engine: the on-device model learns you and the people in your life, the cloud learns nothing
- Attached a receipt to every cloud message showing the exact text that left the device, so the claim is checkable
- Measured it rather than asserting it: an n=100 paired eval cut a frontier attacker's location inference 83% to 3%
- Priced privacy as the product: on-device free and unlimited, metered cloud premium that falls back to device at zero credits
- Proved the free tier makes no network calls with a URLProtocol trap that fails the suite on any request, analytics included

Tosskeep | Asset compiler for game developers

Describe any game asset and get engine-ready output; an AI quality gate decides what is good enough to keep

- Rebuilt its architecture as the art pipeline at Weekend in two weeks: the side project became the day job
- Built meta-evaluation that grades the judges themselves on accuracy and rank correlation against human approve/reject
- Measured that post-processing closes a 28-point quality gap, letting a \$0.03 model match a \$0.36 purpose-trained one
- Anchored the AI judge to images the user already approved, fixing the drift that makes absolute LLM scoring unreliable

animachord | AI character SDK: memory, psychology, social dynamics

Characters remember you, decide whether they trust you, and gossip about each other behind your back

- Grounded the model in published psychology: Big Five traits, attachment theory, Gottman ratios, Ebbinghaus memory decay
- Kept the relationship engine LLM-free: the part deciding whether two characters drift apart has no model call in it
- Modeled gossip as a probabilistic graph walk over trust, with secrets capped at two hops and provenance surviving the trip
- Decayed gossip on two axes, confidence five times faster than intensity, so a rumour stays exciting after it stops being credible
- Split memory into episodic, procedural and semantic stores on separate decay curves, with the identity tier exempt from forgetting
- Reimplemented the whole cognitive engine in Swift rather than binding to it, with CI failing the build if the runtimes diverge on disk

Costa & McCrae (Big Five) - Bowlby, Ainsworth (attachment) - Tulving (memory systems) - Gottman (stability ratios) - Ebbinghaus (retention)

EDUCATION

Lambda School (Bloom Institute of Technology) | iOS Development & Computer Science (career pivot to technical) | 2019-2021

The Wharton School, University of Pennsylvania | B.S. Economics, Finance | 2006-2010

SKILLS & TOOLS

Product: Product Strategy, Roadmapping, 0-to-1, Go-to-Market, Experimentation & A/B Testing, User Research, Pricing, P&L, Hiring

AI/ML: Claude API, Claude Code, Agent Orchestration, Prompt Engineering, Evaluation Methodology, LLM Cost Optimization, Voice AI

Technical: Python, TypeScript, Swift, React, Next.js, FastAPI, Electron, Postgres, MongoDB, Unity, Playwright, GitHub Actions