

> _ OpenPitch

The open, real-time intelligence layer for AI startups —
that any agent can build on.

```
> openpitch get anthropic --arr  
ARR ~$47B [reported · medium] · source: CNBC 2026-05-28 · 🚩 discrepancy flagged
```

A free, open alternative to PitchBook & CB Insights — focused on the AI companies VCs care about.

Private-market data is expensive – and stale.

\$20k–100k

per seat / year

**weeks–
months**

data latency

2–4×

AI-startup growth / yr

Incumbents are accurate because they're human-verified — which makes them slow and costly.

- For a company tripling yearly, a figure verified 6 months ago can be off by multiples.
- The real numbers are already public — founders state ARR on podcasts, filings disclose raises, hiring reveals growth — just scattered, unstructured, and contradictory.
- That's exactly the problem an AI agent is built to solve.

Proof of the problem: our own seed went stale in 5 months.

Company	~Jan 2026	Jun 2026	Δ
OpenAI valuation	\$300B	\$852B	2.8×
Anthropic valuation	\$60B	\$965B	16×
Anthropic ARR	\$4-5B	\$47B	~10×
Cursor valuation	\$9B	\$50B	5.5×
Cognition ARR	\$73M	\$492M	6.7×

Incumbent-grade data goes stale in months.

OpenPitch refreshes it daily, every figure sourced and confidence-scored.

Latency, not coverage — that's the wedge incumbents structurally can't close.

WHAT IT IS

Sourced, confidence-scored intelligence – inside your agent.

1

Every number sourced

Podcast timestamp, filing, or article — no black-box figures.

2

Confidence-scored

Source reliability × speaker × corroboration
× freshness (decays with age).

3

In your agent (MCP/A2A)

Ask Claude Code or Codex directly — no website, no key, no signup.

```
> "What is Cognition's ARR – with sources and confidence?"
```

```
ARR $492M [reported · medium] as of 2026-05 · source: Sacra · +573% YoY
```

How it works – git is the database, runs on free CI.



Daily on free GitHub Actions → commits versioned data → the git log IS the audit trail.

MCP server (local, BYO agent)

Static dashboard

Event feed + A2A agent

The edge nobody else has: it reasons about the numbers.

▮ Contradiction detection

When public sources disagree, it flags it – neutrally.

Mistral valuation:

\$13B confirmed (Series C) vs \$23B rumored
→ public-source discrepancy ▮

Screenshot-worthy, defensible, and unique to OpenPitch.

Σ Derivation / triangulation

Reverse-engineers metrics like an analyst:

ARR = MRR × 12
valuation = round ÷ equity %
ARR = subscribers × ACV
revenue multiple = valuation ÷ ARR

Derived values cross-check reported ones – so fabricated numbers get caught.

Complementary, not a clone – we win a narrow wedge.

	PitchBook / CB Insights	Crunchbase	OpenBB	OpenPitch
Free & open	✗	⦿	✓	✓
Daily freshness	✗	⦿	⦿	✓
In your agent (MCP)	✗	✗	✓	✓
Every figure sourced	⦿	⦿	✗	✓
Contradiction detection	✗	✗	✗	✓
AI-startup focus	⦿	⦿	✗	✓

The honest pitch: the free, fresh, agent-native first look — before you pull the expensive verified report.

TIMING

Why now: agents need grounded data, and the rails just shipped.

97M+

MCP SDK downloads / mo

150+

orgs adopting A2A

~2026

every app gets an AI assistant

The interoperability layer (MCP + A2A) is now the default — and every agent has the same problem:

it hallucinates private-company numbers.

OpenPitch is the free, grounded fact-base agents call so they stop making numbers up — the one place the \$30k incumbents structurally can't follow.

Zero cost – proven, not aspirational.

50

extraction calls/day (demand)

420

free calls/day (capacity)

\$0

marginal cost

How: free GitHub Actions (compute) · git (storage) · free-tier LLM with batching + multi-model rotation · Groq free Whisper · and the user brings their own agent for queries.

Even if we ever paid: ~\$2.70 / month at full 50-company daily scale. The free tier covers it 8× over.

Status: built, tested, and live-validated.

✓ 42 tests passing

Core engine, adapters, extraction, derivation, MCP — all green

✓ MCP connected

Registered in Claude Code (✓) and Codex — answers in-agent

✓ Zero-cost pipeline

Daily GitHub Actions, batching + rotation, \$0

✓ Live-validated

Extracted OpenAI \$852B / \$24B ARR from real sources, matching verified data

✓ Verified seed + 2 contradictions

6 AI companies, real sources, Mistral & Cursor discrepancies

🕒 Remaining

Whisper key, launch-grade contradictions, PyPI publish, demo GIF

Go-to-market: a sharp beachhead, an honest plan.

Beachhead

AI builders who need grounded company data in their own agents/products.

Customer in our channel · recurring, embedded use · the spot incumbents can't reach.

Secondary: founders, VCs/scouts priced out of PitchBook, tech press.

30-day launch

Hook: the contradiction-finder (drama) + free PitchBook-for-AI (utility).

Show HN → X thread → Product Hunt → Reddit, sustained by a renewable drip of new contradictions.

Honest odds of 5,000 stars in 30 days: ~25% (base 2.5-4k).
Stars ≠ PMF — we build for the retained user.

> _

The open intelligence layer any agent can build on.

```
Claude Code:  claude mcp add openpitch -- uvx openpitch-mcp
```

Free · open-source · MCP & A2A native · every number sourced.

github.com/<owner>/openpitch · MIT