JACOB RAMOS

Prompt Engineer — I translate business needs into prompts, and prompts into measured systems

Tel Aviv (relocating September 2026) · +1 813-440-1402 · hire@jacobramos.com · hire.jacobramos.com · spskyway.com · open-civ.com

Founder who runs prompt engineering as a discipline, not an experiment: take a business need, characterize it into a use case with defined success metrics, write the prompt, force the output into a schema, and grade it before a user sees it. Eight months of Open-CIV built that way — five products live in production on multi-agent Claude Code workflows. Before that, four years in revenue seats: the half of this job that means sitting with non-technical people until a requirement is actually a requirement.

EXPERIENCE

Founder — Open Civic Intelligence Ventures (Open-CIV)

Saint Paul, MN · Dec 2025 – Present

- Design and run multi-agent Claude Code workflows as the development process itself: parallel subagents with JSON-Schema-forced structured outputs, phased pipelines, and an adversarial verification stage that has to break the first answer before it ships. MCP servers and custom skills in every repo; the method itself written down — CLAUDE.md agent conventions, reusable skills, workflow scripts — and standardized across six repos so a second person could run it.
- Apply an LLM-as-judge pattern where the judgement costs me something: a structured parallel multi-agent audit of the whole portfolio returning an executive product review — risks ranked P0–P2, per-repo scorecards, every finding independently verified against the code before it was allowed into the report.
- Run GenAI over real unstructured business text: 681 free-text community survey responses collected, the classification and clustering run on the 679 available at export — per-response summaries, named themes with representative quotes, a faceted explorer, and a heatmap that prints its sample size in every cell.
- Define the test scenarios before launch, failures included: a ticketed public event instrumented at 259 QR scan events bucketed by stage — opened, camera granted, decoded, denied, rejected — so a drop-off had somewhere to land.
- Treat output quality as a product surface: provenance labels, data-valid-from floors, minimum-N gating, freshness indicators and explicit degraded states on every number a user sees — anti-hallucination by design. One event that underperformed became a written post-mortem: named failure clusters, P0 fixes, shipped recovery.
- Own the data underneath: Supabase and Postgres schemas and SQL behind a 418-place directory, integrity validations on every load of the ten-tab admin console I run daily, GA4 event instrumentation on every funnel, and a Python on-device counting pipeline — tracking and tripwire logic, SQLite offline buffer, idempotent Postgres sync. Flagship product: 3,100 active users and 38K tracked events in a quarter, organic (GA4, spskyway.com property).

Financial Services Leader → Finance Manager — Automotive Retail

Jan 2021 – May 2025

Morgan Auto Group · Morrie's 394 Hyundai · BMW of Minnetonka — Tampa & Minneapolis

- Owned internet-lead flow as finance manager, training and mentoring 60+ salespeople (Morgan Auto Group): 50–70 closings a month at $500k–$600k volume, $2,800 PVR, 60% attach; earlier, in internet sales: top 3% of 60+ reps, twice Top Salesperson.
- The characterization step before models were in it: turning what a non-technical business said it wanted into a script, a process and a number — daily in the Cox Automotive stack (Dealertrack, VinSolutions, vAuto), downstream of the chatbots that met my shoppers first, which is where you learn what a bad automated answer costs.

Campaigns, Consulting & Client Services

DC / NY / VA / FL · 2016 – 2021

- People leadership and stakeholder translation: staff of 20 across two New York restaurant locations; 200+ volunteers and $50k+ raised as campaign manager; a 79,160-door field program that helped flip a Virginia house seat by 27 votes; a 5k-stakeholder email program at 92% open rate raising $2,000 per send.

SKILLS

AI & LLM: Claude Code (expert, daily since November 2025) — agent orchestration, MCP servers, custom skills, JSON-Schema structured outputs, LLM-as-judge evaluation, prompt and style-system design; n8n, Make and Zapier (working familiarity).

Data & APIs: Postgres, Supabase, SQL, REST APIs and webhooks, edge functions, cron jobs, GA4 instrumentation, Stripe, Resend.

Analysis & delivery: Requirements characterization, use-case and test-scenario definition, success-metric design, written methodology and agent conventions, discovery with non-technical stakeholders, technical documentation.

Engineering: TypeScript/JavaScript, React, Next.js, Swift (native iOS), Vercel, Playwright and Vitest testing.

Languages: English (native), Hebrew (studying now, ulpan from arrival).

EDUCATION

The George Washington University — Elliott School of International Affairs — graduate coursework (2015–2017)