Greg Schwartz

greg@gregschwartz.net
415-547-0392
PromptEngineeringPodcast.com
GregSchwartz.net
github.com/gregschwartz
linkedin.com/in/gregschwartz

SUMMARY

Stanford HCI trained AI engineer with 21 years experience, building exceptional user experiences with **Large Language Models (LLMs)**.

- Ships LLM systems that move operating metrics: 3.2x more patients per clinician, support triage accuracy from 57% to 96%, 6x faster clinical report writing
- Production work with chatbots, evaluation, RAG, vector databases, agents, voice AI and GenUI
- Usage of Anthropic Claude, Claude Code; OpenAI GPT-6, Codex; Google Gemini, Jev, open source LLMs
- Host podcast teaching prompt engineering with **80,000+ downloads**
- Created San Diego's first AI hackathon: 180+ attendees, 18 sponsors
- **Patent US10049598B1**: Passive tracking of food via wearable body sensors (Zipongo)
- Mentor for Springboard (online learning): 96.2% satisfaction from mentees

WORK EXPERIENCE

9am Health (centralized diabetes patient care; HIPAA regulated) Remote
Senior AI Engineer *Jun 2024 - Present*

- Led AI engineering efforts to meet company cost efficiency mandate: as patients grow 2–4x, medical team headcount costs must not grow proportionately. We beat both ends: patients grew 6x (far above goal), while the medical team grew only 1.9x.
- Built HIPAA-compliant LLM patient-support chatbot using OpenAI GPT-4.1, GPT-5-mini, GPT-5.4 + Postgres RAG over 100+ clinical and company protocols customized by insurance company
- Built chatbot that autonomously answers 2/3rds of all patient messages, while message volume increased 8x
- Improved triage accuracy from 57% to **96%** over 36k patient support tickets. Previously needed 3 human support staffers full-time. Iterated on prompts, added + optimized few-shot examples for 15 triage categories, generated golden datasets of 50 and 200 human labeled tickets in Langfuse for eval.
- **Cut token use 25.9%** by optimizing input data, no accuracy regression on golden dataset evals
- Sped up lab results **2.8x** (9.08 min to 3.23 min) replaced human entry with GPT-5.5 parsing lab reports into structured EMR fields. 95% accuracy, with data validation view for clinical review.
- Evaluations via Langfuse, with LLM-as-judge to create golden data set, check prompt changes

MegaPrompt AI (AI service provider) Remote
Fractional CTO & Founder *Aug 2014 - Jun 2024*

Led technical, organizational, and strategic efforts, with an emphasis on 0-to-1 AI strategy, AI engineering, and User Experience that contributed to millions in revenue and startup valuations

- MedowHealth.ai: owned + evaluated 3 prompt chain, **600% speed up** in doctors' report writing
- Sighten.io: rescued solar CRM app, directly leading to **$3.5 million Series A**
- IPVM.com: "Consumer Reports" for security systems, **1200% increase** in users sending invites
- CollegeHunksHaulingJunk.com: **SMS chatbot** answering FAQs, using Twilio and OpenAI GPT-4
- Aurite.ai, brainstorming videos from top Youtube videos: frontend, GPT-4 API calls, prompt chain
- Prompt Engineering IDE, build + evaluate prompts using human feedback: full-stack engineering
- Godaddy: creation, optimization, evaluation of huge prompt library; ads, data analysis, SEO, etc
- California Office of Emergency Services: user research, design solutions

Couple Up (video speed dating app) Remote
Founder, Lead Engineer *Aug 2021 - Feb 2022*

- Built iOS app, backend RESTful APIs, database (Postgres), managed streaming video engineer

Zipongo (healthy eating and meal planning) San Francisco
Director of Product *Sep 2012 - Aug 2014*

- Prototyped, designed, optimized web, iPhone, Android apps; created location-based notifications
- Managed product team (visual design, content, UX design, etc)

Mindswarms (market and user research tool) San Francisco
Lead Engineer & User Experience Designer *May 2010 - Dec 2011*

- **34% growth** in user acquisition following release, with no advertising
- Led and managed 3 Ruby engineers, ensuring seamless collaboration and timely deliveries

Playdom (online games) San Francisco
Front End Engineer & Designer *Sep 2009 - Apr 2010*

- Created in-game notification library and distribution: iPhone (Objective-C), backend (PHP)

Web Advanced (full service website consulting) Irvine
Lead, Full Stack Engineer *Aug 2005 - Jul 2007*

- Built app that manager said was "clearly head and shoulders above what customer or MapQuest was expecting" while learning MapQuest API and new programming language

EDUCATION

Stanford University *Sep 2007 - Jun 2009*
Masters, Human-Computer Interaction, Computer Science
Additional research: haptic feedback in cars, next-generation map interactions

UC Irvine *Aug 2001 - Jun 2005*
Bachelors, Information and Computer Science

SKILLS

Large Language Models	Generative UI	Javascript, Typescript,
Retrieval Augmented Generation	AI Phone Calls: Bland, Retell	React, HTML, (S)CSS
Evals: Langfuse, Braintrust	Text To Speech: ElevenLabs	Postgres, SQL
HCI and Human-AI Interaction	Data visualization, Looker	NoSQL, Redis, KV stores
Function Calling / Tool Calls, MCP	Python, Django	Team leadership

PUBLIC WORK AND AWARDS

Talks: Giving Good Presentations: stop putting paragraphs on your slides; Quantified Dating: analysis of online dating messages, photos, and profiles; Physical activity, calorie expenditure at Burning Man; Are You What You Eat, or What You Buy? Caloric intake, activity, bodyweight, nutrients (became a patent)
Paper: Harvesting Helpfulness, UIST 2008: A/B tested enhancing forum threads with maps
Community: Head Organizer, Quantified Self San Francisco Meetup, 2016 - 2020; Eagle Scout, 1999; Fire dancer, poi (you've probably seen it in Hawaii), 2010
Awards: iOSDevCamp 2011: Best Audio/Video App; iPhoneDevCamp 2008: Best Game; UCI Graduate School of Management Business Plan Competition: Best Undergrad Team and Third Place: 2+ years before the Kindle release, proposed eBook reader and integrated content delivery system

FEEDBACK

Greg completed the most incredible prompt I've ever seen, having spent a ton of time myself researching it and not finding any examples anywhere. Greg was very aware of the needs of our audience, adding elements to his prompts to ensure output was optimized, reducing the fluff that tends to come with Chat GPT. This is our second time working with Greg and he will continue to be our go-to prompt engineer.
— Alycia Leno, GoDaddy.com, Marketing Director

Greg is a highly skilled AI automation engineer with a deep subject matter expertise. It was great to have his advice and support at this pivotal stage of our journey. Would highly recommend him.
— Upwork Client, Project: AI for automated reports from medical transcripts

Huge asset to the team. Appreciated his ability to tackle and present his work from the perspective of user needs and real use cases.
— Genevieve Wang, Zipongo, Product Head