

Brian Mar

Forward Deployed Engineer

(510) 717-1898 | brian099121@gmail.com | linkedin.com/in/brnmr | github.com/moujiandao | Bay Area, CA

PROFESSIONAL SUMMARY

Enterprise deployment engineer with 7+ years shipping production systems for 30+ mid-to-enterprise clients - full ownership from technical discovery through production handoff, across data infrastructure, CRM, and marketing analytics. Now applying that track record to AI: multi-agent orchestration, LLM-powered automation, agentic workflow design, and production-grade eval pipelines. Built for the FDE role - fast prototyping, client-facing technical delivery, and working systems over polished demos.

CORE COMPETENCIES

Client-Facing Technical Discovery	Enterprise Deployment (30+ Clients)	Multi-Agent Orchestration	LLM API Integration & RAG Pipelines
Agentic Workflow Design	Eval Engineering & LLM-as-Judge	dbt / Snowflake / Airflow / Python	Prototype-to-Production Delivery

WORK EXPERIENCE

B&E Medical LLC

June 2024 - May 2026

Agentic AI Engineer -- Union City, CA

- Designed LLM-as-judge eval harness to enforce brand voice and style compliance in AI-generated product descriptions; rubric dimensions derived from medical supply domain expertise, covering terminology accuracy, structural format, and regulatory tone; failure-mode categorization drove targeted prompt refinements across iterations.
- Deployed end-to-end advertising automation system leveraging Amazon Seller Central API and Sponsored Products Bulk Operations to operationalize campaign optimization, automating 90% of bid adjustment decisions.
- Built AI-assisted product listing pipeline using LLM APIs to auto-generate structured product descriptions from data scaffolds, enabling a rapid review-and-publish workflow.
- Implemented data-driven bid optimization logic that identifies underperforming campaigns and reallocates budget toward highest-converting keywords.

Full Circle Insights

Feb 2019 - May 2024

Solution Architect -- San Mateo, CA

- Deployed marketing attribution data models tracking \$100M+ in spend across multi-touch campaigns for 30+ mid-to-enterprise clients -- owned full implementation from technical discovery and scoping through production handoff.
- Built Python pipeline resolving a core attribution gap: joined anonymous web visitor data with identified CRM records to surface full-funnel visibility from first digital touch to closed-won revenue, driving measurable reallocation of marketing budgets toward highest-converting channels.
- Led technical discovery and solution scoping with marketing and revenue operations leaders, translating ambiguous business requirements into scalable data models and funnel analytics frameworks.
- Established sandbox-to-production migration protocols preventing data loss through staged validation and quality checks across 20+ enterprise implementations.
- Ran data audits across CRM, marketing automation, and analytics platforms -- surfaced pipeline bottlenecks and data inconsistencies before they impacted client reporting.

Full Circle Insights

Dec 2017 - Feb 2019

Senior Business Analyst -- San Mateo, CA

- Led end-to-end technical implementations across CRM and analytics platforms -- translated business requirements into data models, trained client stakeholders, and ran regular data quality audits to surface integrity issues before they impacted reporting.

Salesforce Data Analyst – San Francisco, CA

- Built automated workflows and Salesforce data integrations under strict enterprise security constraints (no direct DB access, VPN-only API surface, secret vault), reducing manual data entry by 70-80% and improving data consistency across the US West Salesforce instance.

PROJECTS

Career-Ops Multi-Agent System 2026 - Present

Claude Code skill orchestrator turning raw job postings into tailored applications across the full lifecycle: discovery, evaluation, autofill, interview prep, and tracking. Zero-LLM-cost portal scanner hits Greenhouse / Ashby / Lever APIs directly, feeding an Anthropic-powered evaluation pipeline that scores fit on a 7-block rubric and graduates matches into a SQLite store with Datasette dashboards. Playwright-driven autofill across 30+ portals handles React-controlled inputs, cascading dropdowns, and cross-platform variations under human-in-the-loop review. Eval rubric scores AI-generated reports for keyword coverage and reasoning quality; pipeline integrity tests run as CI quality gates. Persistent cross-session memory via claude-bridge retains agent context across runs.

Python, Anthropic API, Claude Code Skills, Playwright, SQLite, Datasette, Multi-Agent Orchestration, Eval Engineering, CI Quality Gates

DinkGPT - Pickleball AI Playbook 2025 - Present

Production AI app delivering optimal shot selection and court positioning for every in-game situation. Custom eval harness with multi-dimensional LLM-as-judge rubric scores strategic soundness, reasoning quality, and specificity. Failure-quadrant error analysis surfaces whether regressions trace to prompt, retrieval, or judge calibration; CI gate enforces minimum pass rates on the golden set before any change merges. RAG over a curated pickleball knowledge base; human-in-the-loop feedback from beta users continuously refines recommendations.

Python, Anthropic API, RAG, Vector Search, LLM-as-Judge, Eval Engineering, Error Analysis, Human-in-the-Loop Training

SwingCoach - AI Tennis Coaching Pipeline 2025

Production pipeline turning user swing video into structured Claude-generated coaching feedback. Dual-track eval system: 9 deterministic structural checks run inline on every generation, plus an offline LLM-as-judge rubric for qualitative assessment of coaching specificity and actionability. Self-repair loop retries with targeted prompt adjustments when feedback fails structural checks, raising feedback acceptance rate without inflating production latency.

Python, Claude API, LLM-as-Judge, Custom Evals, Self-Repair Loops, Prompt Engineering

EDUCATION

Bachelor of Science, Mechanical Engineering - Minor: Computer Science

University of California, Davis

CERTIFICATIONS

Salesforce - 14x Certified (Data Architect track) - salesforce.com/trailblazer/brianmar

Microsoft - DP-900 Azure Data Fundamentals - Credential ID: D7BD1D0E0C1055C7

SKILLS

AI & Agents: LLM APIs (OpenAI, Anthropic), Multi-Agent Orchestration, RAG Pipelines, Prompt Engineering, Eval Engineering, LLM-as-Judge, Human-in-the-Loop Training, LLM Fine-Tuning, LLM Observability (Langfuse, Braintrust)

Data Engineering: dbt, Snowflake, Airflow, Databricks, Redshift, S3, Dimensional Modeling, ELT/ETL **Programming:** Python, SQL, Pandas, Git

Cloud: AWS (S3, Glue, Lambda, Athena, EC2), Azure (Data Lake Gen2, Data Factory, Synapse Analytics)

CRM & Integrations: Salesforce, Marketo, HubSpot, Marketing Cloud, REST APIs

Automation: Playwright, Agentic Workflow Design, CI Quality Gates