

Mihir Parmar

Founder & Product Leader | Luxury Travel · AI-Operated Advisory

+1 (908) 745-9085 | parmarmihir@gmail.com | San Francisco Bay Area
mihirparmar.ai | linkedin.com/in/mihir-r-parmar | github.com/mihirparmar-ai

SUMMARY & SKILLS

Founder and product leader at the intersection of travel and AI - the ICP travel-tech builds for: a working luxury travel advisor (**\$700K+ in bookings in 2026 to date**, 100% repeat/referral growth) who runs the advisory on an **AI operating stack built from scratch** - Claude Code agents + skills, **the agent that drafts every client itinerary** with its own eval loop and guardrails, MCP servers - covering research, itinerary assembly and QA, client lifecycle, and content. Behind that: a decade-plus of product leadership on high-trust platforms - Zillow's **rentals marketplace** (digital leases: **doubled landlord activation**, share **18% to 30%**) and **document intelligence (1M+ docs/year)** - plus banker workflow platforms at Goldman Sachs, Deutsche Bank, Credit Suisse, and UBS. I know travel from the inside - suppliers, margins, seasonality, what travelers actually need - and I build products that scale that judgment.

Core Skills & Toolkit:

- **Travel & Hospitality:** Trip design & itinerary QA • Supplier/DMC sourcing & negotiation • Pricing & margin management • Client lifecycle & retention • Destination intelligence (190-destination corpus)
- **AI / Data Systems:** Agents that draft and check itineraries (Claude Code skills, MCP servers) • Evals & AI quality (golden sets, precision/recall, dispositions log) • Failure-mode catalogues • LLM product design • Human-in-the-loop + guardrails (act/draft/escalate) • Experimentation / A/B testing • SQL analysis
- **Product Leadership:** 0 to 1 + scale • Marketplaces & two-sided growth • Customer discovery & user research • GTM & launches • Financial modeling / ROI
- **Domains:** Travel & Leisure • Marketplaces • PropTech • FinTech • Regulated Workflows
- **Hands-on tools:** Claude Code (custom skills, hooks, MCP servers, scheduled tasks) • Claude • ChatGPT • Gemini • Supabase • Vercel • GitHub Actions • SQL • Python/Node scripts • Replit • Lovable • Figma Make • ElevenLabs

PROFESSIONAL EXPERIENCE

En Route Luxe - Founder / Product Lead (solo, AI-operated luxury travel advisory) | SF Bay Area | 2025 - Present:

Boutique luxury travel advisory built on the idea that luxury should feel personal, not transactional, and run as the production environment for an agentic system I built and operate alone.

- **Business: \$700K+ in bookings in 2026 to date**, grown entirely on repeat and referral.
- **Discovery is the job, not a phase:** every client starts with an intake call I run; **20+ families a year, 50+ trips, 30+ countries**, each itinerary built from that call and the family's past trips.
- **End-to-end product:** destination intelligence, curated itinerary design, **supplier and DMC sourcing and negotiation across four continents**, pricing and margin management, payment operations.
- **The product work itself runs on an agent I built:** the itinerary builder compiles a client profile (explicit intake answers, conversational signals from the call, behavioral signals from past trips) into constraints, drafts the day-by-day itinerary, builds the client page, and runs its own check step before anything reaches me. Around it: custom **Claude Code skills** for recurring workflows (itinerary QA, research ingestion, post editors), a flight-search **MCP server**, a 190-destination corpus, and the client site and itinerary pages I prototyped myself (Lovable, then Claude Code) - hands-on with tool use, MCP, and the real failure modes of agents in production.
- **Evals, guardrails and human-in-the-loop by design:** a golden set of trips with expected findings, a checker scored on recall and precision, a scorecard per draft, and a dispositions log where each finding I accept, reject or fix feeds per-class precision. Seven failure classes catalogued from production and turned into checks: closures on the day scheduled, ticket windows on a clock (Shibuya Sky sunset slots release 14 days out and sold out in about 12 hours), transfer buffers, contingency paths, jet-lag pacing, verified versus estimated pricing, stale venues. The critic runs read-only in a fresh context; agents draft, the human reviews, as at Zillow.

- **Deliberate automation boundaries, enforced in tooling:** every action class is act alone, draft, or escalate; the model can draft an email but the Gmail permission is create-draft only; every site change goes through a branch, a pull request, and my merge. Judgment calls - supplier selection, disruption response, taste - stay human, and several workflows stay deliberately simple because an agent would not have earned its keep. **Knowing where AI adds real leverage and where it does not** is the product thesis and the subject of my public writing.

Zillow Group - Principal Product Manager, Documents + Rentals Marketplace | SF Bay Area | 2019 - 2026:

Built Zillow's document intelligence platform end to end and helped grow the rentals marketplace.

- **Zillow Rentals (flagship proof at scale):** shipped digital leases nationally in ~12 months; instrumented and analyzed the create, send and sign funnel in SQL, and iterated with experiments - **doubled landlord activation** and grew **market share from 18% to 30%**.
- **Shipped multimodal LLM document understanding to production and chose the gate that made it trustworthy.** Extraction was one platform use case; OCR and NLP topped out near 85% and the expected fix was a better OCR vendor. Average accuracy was the wrong target, because a contract that is 92% right is unusable if nobody can tell you which 8% is wrong, so we gated on **per-field confidence** instead, with **human-in-the-loop guardrails** routing low-confidence fields to review. Accuracy went from ~85% to **98%+**, with **~100% on reviewed compliance-critical fields**, and the automation absorbed **~\$2M a year** of manual labeling and review.
- **Built the multi-tenant Document Platform first, then deployed it team by team:** discovery mapped the use cases every business ran through DocuSign and dotloop; we brought them in-house as foundational services (e-sign, storage, compliance, viewing, sharing, fillable forms) and rolled the platform into **every Zillow business and internal team that touched a document** (Rentals, Home Loans, Zillow Offers, Zillow Closing Services, Premier Agent, plus marketing and sales, legal, recruiting, vendor contracts; **20+ teams**), each a deployment and launch with its own Director/VP stakeholder, baseline, success metric, and GTM and adoption playbook: **1M+ docs/year** at **95%+ platform accuracy**, **5+ use cases/quarter** on the APIs, **~\$2M/yr** e-sign licensing retired.
- **Owned model economics: ~\$7M annual budget;** led build-vs-buy across 3 vendors + internal PoC, scored per field on shared golden sets with cost/latency in the bake-off - integrated in ~4 weeks vs a ~6-month build, avoiding **~\$2M/yr** of internal platform engineering - and stayed model-agnostic.

Deutsche Bank, Credit Suisse, UBS | Product Management (banker and prime brokerage platforms) | New York & Singapore | 2013 - 2019:

- Led a **25-person global team** (engineering, design, data) that shipped **iOS and web workflow tools and analytics** for investment bankers, raising productivity **~30%** where nothing ships without compliance sign-off.
- Won the **\$5M** prime brokerage re-engineering case: **10+ clients**, **\$50M annual revenue**, **\$500K/yr** cost cut.

Deloitte | Management Consultant | New York | 2011 - 2013:

- **100+ metric vendor scorecard** that settled **buy-versus-build** and roadmap for a hedge fund's portfolio platform.

Goldman Sachs | Software Engineer | New York | 2007 - 2011:

- Engineer on a global portfolio and trading platform (**\$100B+ AUM**, **10K+ daily transactions**); built trade-booking automation that cut manual errors and **saved \$600K+ a year**.

EDUCATION

New Jersey Institute of Technology - M.S., Computer Science	2006 - 2007
Vishwakarma Institute of Technology (University of Pune) - B.E., Electronics and Telecommunication	2001 - 2005