

Parik Ahlawat

Senior AI Consultant & Engineer | AI Agents, Automations & Full-Stack AI Systems

pa.parikahlawat@gmail.com | Remote | [linkedin.com/in/parik-ahlawat](https://www.linkedin.com/in/parik-ahlawat) | github.com/par1kahl | [par1kahl](#)

PROFESSIONAL SUMMARY

Senior AI Consultant and Engineer with 6+ years architecting and building production AI systems for businesses. 486 automations shipped across 22 projects spanning 10 industries. I help companies go from "we want AI" to systems that run in production, generate revenue, and scale without breaking.

Proven track record: improved LLM accuracy 15-25% on industry benchmarks, reduced developer effort 60% through AI-powered workflows, and built a 77-automation ad creative platform processing thousands of assets for a marketing agency. Specialize in multi-agent orchestration, RAG pipelines, full-stack AI applications, and business process automation.

CORE COMPETENCIES

Consulting & Strategy	Engineering & Architecture
AI system architecture reviews	Multi-agent orchestration (66 agents built)
Production readiness audits	RAG pipeline design (27 pipelines shipped)
Cost optimization and model routing	Full-stack AI apps (FastAPI, Next.js, PostgreSQL)
Requirements to production planning	LLM evaluation frameworks (BFCL, TAU, LiveBench)
AI integration into existing stacks	n8n and workflow automation (4,274 workflows)

TECHNICAL SKILLS

AI & Machine Learning: LLM Engineering (OpenAI, Claude, Gemini, Mistral, Llama) | Multi-Agent Orchestration | RAG Pipelines | Prompt Engineering | Model Routing & Cost Optimization | LangChain, LangGraph | Neurosymbolic Reasoning | Small Reasoning Models (SRMs) | Evaluation Frameworks (Gorilla BFCL, TAU-bench, LiveBench) | MCP Server Development

Software & Backend: Python | TypeScript | FastAPI | Node.js | Express | Hono | REST APIs | GraphQL | Celery | Bull

Full-Stack & Frontend: React | Next.js | Tailwind CSS | Three.js | Vite | Recharts | XYFlow

Data & Storage: PostgreSQL | MongoDB | Redis | SQLite | Qdrant | Pinecone | ChromaDB | Neo4j | pgvector | Supabase

Automation & Integration: n8n | Slack API | Telegram Bot API | Meta Ads API | Google Sheets API | Dropbox API | Stripe | Apify | fal.ai | ElevenLabs

DevOps & Infrastructure: Docker | Docker Compose | Kubernetes | AWS | GCP | GitHub Actions | CI/CD | PM2 | nginx

PROFESSIONAL EXPERIENCE

Senior AI Consultant and Engineer | Independent

Jan 2019 – Present | Remote

Architect and deliver production AI systems for businesses across FinTech, advertising, SaaS, logistics, media, and developer tools. 486 automations, 22 projects, 147 unique technologies.

Client: Marketing Agency (Advertising/Marketing)

- Built a 77-automation AI ad creative platform with 5 sub-systems: competitor intelligence scraping, template-based ad generation, native ad generation, hook ad generation, and video ad pipeline with avatars, B-roll RAG, and scene workbench
- Integrated Meta Ads API for automated publishing, Slack approval workflows, and Google Sheets tracking
- Multi-LLM orchestration (Claude, Gemini, OpenAI, fal.ai, ElevenLabs) with Bull/Celery job queues

Stack: TypeScript, Python, FastAPI, Express, Next.js, PostgreSQL, pgvector, Redis, Celery, Bull, Meta Ads API

Client: Sales/CRM Platform

- Built AI sales agent using LangGraph with agentic RAG and Text-to-SQL for natural language lead nurturing
- Multi-node agent graph: router, RAG retrieval (ChromaDB), Text-to-SQL (Vanna), response generation

Stack: Python, LangGraph, LangChain, OpenAI, ChromaDB, Django Ninja, PostgreSQL

Client: FinTech/Trading Platform

- Built visual node-based strategy builder (XYFlow) with backtesting engine and real-time market data via WebSocket
- Zerodha broker integration, payment system, 2FA auth, Sentry monitoring, CI/CD with Trivy security scanning

Stack: TypeScript, Python, Next.js, FastAPI, PostgreSQL, Celery, Zustand

Client: Transport/Logistics (Netherlands)

- Built AI zero-emission zone intelligence app: ROI calculator, exemption tracker, corridor analyzer, AI chat assistant

Client: Marketing/SaaS

- Built lead qualification quiz platform with scoring engine, visitor tracking, admin dashboard with analytics

Founding Senior AI Engineer | CoreThink AI

Mar 2025 – Aug 2025 | Remote

First engineering hire. Built the LLM evaluation and reasoning infrastructure from scratch. 30 automations.

- Designed unified evaluation framework integrating Gorilla BFCL (function calling), TAU-bench (multi-step reasoning), and LiveBench (dynamic evaluation) with common base_url abstraction for parallel multi-model testing
- Built neurosymbolic Chain-of-Thought pipeline: NL processing, symbolic representation, inference engine, and 4-layer verifier (structural, logical, cross-check, failure detection)
- Created metrics system: accuracy, step correctness, tool precision, schema validity, reasoning depth, consistency, verifier agreement scores
- **Impact: LLM accuracy +15-25% on 4 benchmarks. Agent latency -30% via batched vector search and caching.**

Product Manager and Senior AI Engineer | CraftifAI

Oct 2025 – Dec 2025 | Remote

Led AI agent product for embedded and edge platforms where cloud-based models are not an option.

- Designed agent workflows operating within real hardware constraints (memory, latency, connectivity)
- Built model orchestration layer selecting between on-device and cloud inference based on resource availability
- Managed product roadmap and engineering execution simultaneously

AI Engineer and Researcher | AccelChain

Apr 2023 – Aug 2025 | Remote

Built AI-powered developer tools for automated code analysis, generation, and security auditing. 65 automations.

- Built multi-stage AI code generation pipeline: task normalization, RAG-style context injection, multi-pass LLM generation, validation layers, iterative refinement loops
- Designed 4-agent coding workflow: planner, coder, reviewer, fixer with controlled iteration loops and cross-agent state
- Built automated audit agents with self-critique loops: static analyzer, LLM reasoning, structured feedback generator
- Created Solidity AST parser with security scanning engine (5 scan types: call, external, operator, transfer, constructor)
- **Impact: Developer effort -60%. Two foundation grants secured for AI tooling research.**

Freelance Engineer | Sidechain Labs

2019 – 2023 | Remote

Backend systems, APIs, AI tooling, and blockchain integrations for startups and enterprises.

- Delivered 20+ client solutions across backend architecture, API design, and open-source AI tooling
- Collaborated with engineering teams at Google and other major tech companies
- Published open-source agents and reasoning modules used by community developers

KEY PROJECTS

Kronus | Open-Source Agent Orchestration Framework kronus.tech

108 automations. 21 specialized agents, 27 skills, MCP server integrations, Telegram control plane daemon, knowledge graph brain (13 tools), React dashboard, hub server with JWT auth and Stripe billing. 210 commits.

Tech: TypeScript, Bun, React, Vite, Hono, PostgreSQL, Redis, Stripe, Grammy, D3.js, MCP SDK

CommsOS | Meeting Intelligence Platform

23 automations. RAG over 475+ meeting transcripts with Qdrant vector search. Intent-aware search, AI chat with streaming, action item extraction with Slack delivery, executive summaries, people profiles.

Tech: Python, TypeScript, FastAPI, Next.js, Qdrant, OpenAI, Gemini, PostgreSQL, Redis, Celery

Meta Ads MCP Server | 34 Tools

34 MCP tools for Meta ad management and competitor intelligence. 26 marketing tools and 8 competitor tools. Apify ingestion pipeline for competitor data.

Tech: TypeScript, MCP SDK, PostgreSQL, Apify, Zod, Vitest

Phavella | AI Business Assessment Platform phavella.com

18 automations. Three.js 3D hero, AI business assessment wizard generating personalized plans with ROI projections and roadmaps, email campaign system, admin dashboard.

Tech: TypeScript, Python, Next.js, FastAPI, Three.js, GSAP, Recharts

KEY ACHIEVEMENTS

- **486 automations** shipped across 22 projects in 10 industries
- **66 AI agents** built for production use cases
- **27 RAG pipelines** designed and deployed
- **34 MCP tools** in a single Meta Ads server
- **15-25% LLM accuracy improvement** on 4 benchmarks
- **60% reduction** in developer effort via AI workflows
- **30% latency reduction** via batched vector search
- **100% Job Success** on Upwork
- **Two foundation grants** for AI tooling research
- Creator of **Kronus**, open-source agent framework

EDUCATION

B.S. Computer Science | Purdue University, West Lafayette, IN, USA | 2018