

Daniel Parker

Senior AI Software Engineer

danielkparker94@gmail.com | (435) 276-0407 | Naperville, IL | LinkedIn

SUMMARY

Senior AI Software Engineer with 8+ years of experience building custom software, e-commerce platforms, machine-learning systems, and production AI applications. Experienced in architecting end-to-end AI agents, voice automation, retrieval-augmented generation, semantic search, recommendation systems, CRM integrations, and customer-facing SaaS products using Python, FastAPI, TypeScript, React, LangGraph, OpenAI, Vapi, Deepgram, PostgreSQL, and AWS. Proven ability to take complex products from initial requirements and system design through development, evaluation, deployment, observability, and continuous improvement. Combines strong full-stack engineering fundamentals with practical AI implementation, product judgment, technical leadership, and measurable business impact.

SKILLS

Full-Stack Engineering — Python, TypeScript, JavaScript, SQL, React, Next.js, Redux, Tailwind CSS, FastAPI, Node.js, Express.js, Django, REST APIs, GraphQL, WebSockets, microservices, event-driven architecture

Generative AI, Agents & Retrieval — OpenAI API, Anthropic Claude, Amazon Bedrock, OpenAI Agents SDK, LangGraph, LangChain, retrieval-augmented generation, function calling, MCP, structured outputs, prompt and context engineering, agent orchestration, embeddings, semantic search, hybrid retrieval, metadata filtering, reranking, model evaluation, fine-tuning

Conversational & Voice AI — Vapi, Deepgram, ElevenLabs, Amazon Transcribe, Amazon Polly, Twilio Programmable Voice, SIP, WebRTC, streaming transcription, text-to-speech, voice activity detection, endpointing, interruption handling

Machine Learning — PyTorch, TensorFlow, XGBoost, scikit-learn, pandas, NumPy, Amazon SageMaker, Amazon Personalize, recommendation systems, collaborative filtering, learning-to-rank, classification, time-series forecasting

Data & Search Platforms — PostgreSQL, pgvector, Redis, Pinecone, Elasticsearch, Amazon OpenSearch Service, MongoDB, MySQL

Automation & Integrations — n8n, Make, GoHighLevel, HubSpot, Salesforce, Google Ads API, Meta Marketing API, webhooks, CRM automation, appointment scheduling

Cloud, DevOps & MLOps — AWS, Docker, Kubernetes, Terraform, GitHub Actions, CI/CD, ECS, Lambda, SQS, EventBridge, Step Functions, RDS, S3, Airflow, MLflow

Observability, Security & Quality — CloudWatch, OpenTelemetry, Langfuse, Sentry, automated testing, OAuth 2.0, RBAC, secrets management, structured logging, audit trails

PROFESSIONAL EXPERIENCE

Proven ROI, Senior AI Software Engineer

09/2023 – 08/2026 | Austin, TX

- Led the architecture and delivery of a multi-tenant AI revenue-operations platform built with React, Next.js, FastAPI, LangGraph, PostgreSQL, pgvector, Redis, and AWS, supporting 18 client workspaces and 75+ configurable AI and automation workflows.
- Built inbound and outbound voice agents for lead qualification, scheduling, customer support, campaign follow-up, and after-hours coverage, processing approximately 4,500 calls per month using Vapi, Deepgram, OpenAI, Anthropic, ElevenLabs, and Twilio.
- Introduced a provider-independent speech and model layer across Deepgram, Amazon Transcribe, ElevenLabs, Amazon Polly, OpenAI, Anthropic, and Bedrock, allowing services to route by latency, quality, availability, and cost while failing over automatically during disruptions.
- Connected voice and chat agents with HubSpot, Salesforce, and GoHighLevel, giving agents access to customer context and enabling lead scoring, CRM updates, meeting scheduling, confirmations, and qualified-call transfers.
- Improved grounded-answer accuracy by 24% by developing permission-aware RAG services over websites, pricing data, policies, campaign history, and sales documentation, combining pgvector, OpenSearch, metadata filtering, reranking, citations, and confidence-based escalation.
- Reworked lead-management automation around EventBridge, SQS, Lambda, n8n, idempotent APIs, retry policies, audit logs, and approval gates, reducing repetitive work by 61% and average response time from 26 minutes to under three minutes.
- Standardized deployment, observability, and AI evaluation with Docker, ECS Fargate, Terraform, GitHub Actions, CloudWatch, OpenTelemetry, Langfuse, and Sentry, creating an evaluation suite with 450+ test conversations, adding safeguards for high-risk agent behavior, and mentoring four engineers.

Faire, AI/ML Software Engineer

01/2021 – 08/2023 | San Francisco, CA

- Delivered marketplace capabilities for supplier onboarding, catalog management, wholesale ordering, retailer accounts, payments, inventory visibility, and order tracking using React, TypeScript, FastAPI, PostgreSQL, Redis, and AWS, supporting approximately 9,000 retailers, 1,200 suppliers, and 620,000 product listings.
- Designed personalized recommendation services with collaborative filtering, PyTorch ranking models, product embeddings, behavioral data, and Amazon Personalize experiments, increasing recommendation click-through rate by 15% and average order value by 7%.
- Rebuilt product discovery around Elasticsearch, Amazon OpenSearch Service, synonym expansion, behavioral ranking signals, learning-to-rank models, Pinecone, and OpenAI embeddings, reducing zero-result searches by 22%.
- Automated product classification, attribute extraction, supplier-data normalization, duplicate detection, and description generation across approximately 140,000 monthly catalog updates, cutting manual review effort by 46%.
- Trained XGBoost and neural-network forecasting models using order history, seasonality, promotions, pricing, supplier lead times, and category trends, reducing forecast error by 14% and stockouts by 17%.
- Productionized ML and generative-AI services through SageMaker, FastAPI, ECS, Lambda, Airflow, Step Functions, Docker, MLflow, and CloudWatch, maintaining 99.9% availability, improving inference latency by 36%, and increasing schema-valid outputs from 71% to 88% through a GPT-3.5 Turbo support pilot.

KitelyTech, Full-Stack Software Engineer

06/2017 – 12/2020 | Chicago, IL

- Shipped 11 responsive web applications and SaaS products using React, TypeScript, Node.js, Python, Django, PostgreSQL, and AWS, enabling clients in professional services, retail, logistics, and business operations to improve workflow efficiency.
- Created 35+ reusable REST API endpoints for authentication, account management, payments, scheduling, document processing, notifications, and reporting, integrating Stripe, Salesforce, Twilio, Google Maps, and logistics platforms to accelerate client integration times and reduce development effort.
- Developed a multi-tenant operations-management platform with role-based access control, configurable workflows, reporting dashboards, audit histories, and customer self-service capabilities used by more than 2,400 business users.
- Modernized four legacy applications into modular cloud-hosted services using Docker, AWS EC2, RDS, S3, Nginx, automated database migrations, and GitHub Actions, increasing deployment frequency from biweekly to three or four releases per week.
- Strengthened performance and reliability through Jest and Pytest automation, SQL query optimization, caching, structured logging, monitoring, and root-cause analysis, reducing average API response time by 41% and production incidents by 29%.

EDUCATION

University of Chicago, Bachelor's degree, Computer Science

2013 – 2017 | Chicago, IL