

MATT KILMER

AI Product Manager & Software Engineer

Brooklyn, NY | mvkilmer@gmail.com | 917-327-7141
mattkilmer.com | github.com/MattKilmer | linkedin.com/in/matt-kilmer

SUMMARY

AI product manager and software engineer specializing in production LLM systems. I ship code and build 0→1 products - from AI-powered SaaS tools processing thousands of items at scale to autonomous development workflows. Deep technical understanding of LLM constraints: context window optimization, cost-per-request economics, handling non-deterministic outputs, and building evaluation frameworks for subjective quality. 10 years product management experience across technology, entertainment, and consumer goods.

TECHNICAL SKILLS

AI/ML Product:	LLM Integration & API Design • Prompt Engineering & System Design • Context Window Optimization • Structured Outputs & Function Calling • RAG System Architecture • Prompt Caching & Cost Optimization • Evaluation Frameworks • Safety & Content Filtering
Engineering:	TypeScript/JavaScript • React/Next.js • Node.js • Python • PostgreSQL • API Design • Git/GitHub • AWS/Vercel
Product:	AI Product Strategy • 0→1 Product Development • Roadmap Planning • User Research • Stakeholder Management • Go-to-Market • Agile/Scrum • Product Analytics

PROFESSIONAL EXPERIENCE

Viewcy, New York, NY Senior Product Manager	June 2021 - Present
---	---------------------

- Oversee core product roadmap for event ticketing and live streaming platform serving artists, fans, and sponsors
- Partner with executive stakeholders to define roadmap and set goals aligned with company vision and mission
- Led evaluation and integration of LLM capabilities: real-time content moderation, automated highlight generation, natural language search across video archives
- Own product specification and development process, interfacing with engineering, design, UX, marketing, and analytics teams to consistently deliver releases on time
- Balance needs of multiple user classes to drive success for all customers and stakeholders

Doji Natural Products, New York, NY Founder and CEO	August 2019 - Present
---	-----------------------

- Oversee all aspects of business operations and strategy for boutique eco-conscious health and beauty brand
- Establish and maintain relationships with vendors and strategic partnerships
- Manage P&L, supply chain, e-commerce operations, and go-to-market strategy

Jammer, New York, NY
Senior Product Manager

January 2017 - February 2019

- Led product development lifecycle for flagship music app and interactive digital music service
- Measured and evaluated product optimizations for business impact: conversion, retention, LTV
- Hired 10+ employees across product, engineering, sales, and design
- Secured partnerships with Scooter Braun, Universal Music Group, Warner Music Group
- Shipped features used by 100K+ active users with measurable engagement improvements

Keezy, New York, NY
Product Manager

September 2015 - January 2017

- Set product vision and strategy aligned with business objectives for music software startup
- Defined releases, oversaw user testing, prioritized features, owned product roadmap
- Led development of three apps with 1M+ downloads across iOS and Android
- Collaborated with design and engineering to ship high-quality, performant mobile experiences

FX Networks, New York, NY
Music Coordinator

March 2010 - September 2015

- Oversaw all aspects of music production for Emmy Award-winning television show (5 seasons)
- Interfaced with creative team, hired 50+ musicians, managed asset delivery on tight schedules
- Coordinated complex productions involving multiple stakeholders and hard deadlines

AI PROJECTS (BUILT 0→1)

Claude Slackbot | Autonomous Development Assistant

TypeScript • Node.js • Anthropic Claude API • Slack API • GitHub API • Vercel • Function Calling

Production bot that analyzes codebases, generates fixes, creates pull requests, and deploys previews - all from natural language Slack messages. Built autonomous Slack bot powered by Claude that maintains conversation context across multiple interactions, analyzes entire codebases using semantic search, and provides real-time status updates. Designed system prompts to maintain code quality standards and handle edge cases gracefully.

Key Technical Challenges:

- Context management through thread-aware memory and codebase summarization to handle large repos within token limits
- Function calling integration for GitHub/Vercel APIs with proper error handling
- Human oversight points: bot proposes changes but requires approval before merging
- Cost optimization through prompt caching and selective API calls for production viability

BatchSEO for Shopify | Commercial AI SaaS Platform

Next.js • TypeScript • Shopify API • LLM Integration • PostgreSQL • Vercel • Structured Outputs • Prompt Caching

Production SaaS helping Shopify merchants optimize product listings at scale through intelligent content generation and bulk processing. Built Shopify app with streaming LLM integration for real-time content generation preview, structured output validation for consistent formatting, bulk job processing with progress tracking and failure recovery, and prompt engineering system that incorporates brand voice examples and product metadata.

Key Technical Challenges:

- Quality control workflows (preview-before-apply) to prevent production errors and build merchant trust
- Validation layers checking for hallucinated product details, maintaining keyword density targets, and preserving brand voice
- Cost-per-request optimization through prompt caching, batched API calls, and smart context window management
- PostgreSQL for job queue management and rollback history to enable error recovery

TradingGPT | Natural Language to Executable Strategies

Next.js • TypeScript • Python • LLM Integration • Financial APIs • PostgreSQL • Function Calling • Structured Outputs

Platform translating plain English into validated, backtested algorithmic trading strategies with risk management and paper trading deployment. Built natural language processing pipeline that parses strategy descriptions, validates against financial logic constraints, generates Python execution code using function calling, runs historical backtests with transaction cost modeling, and visualizes performance metrics and risk profiles.

Key Technical Challenges:

- Financial logic validators that catch strategy errors LLMs miss: survivorship bias, transaction costs, realistic slippage, and position sizing
- Visualizing backtest results (equity curves, drawdown charts, correlation matrices) across market conditions
- Safety constraints and risk limits enforced at execution layer, not just generation layer
- Structured outputs ensuring generated code follows safety standards while remaining financially sound

EDUCATION

Berklee College of Music - Boston, MA

Bachelor of Arts in Music Business Management