

Wyatt Walsh

SENIOR AI/ML ENGINEER

wyattowalsh@gmail.com (209) 602-2545 New York City, New York LinkedIn: wyattowalsh GitHub: wyattowalsh

Senior AI/ML engineer blending agentic AI product work, data-platform depth, and product-minded full-stack execution across fintech, enterprise, and open-source environments.

EXPERIENCE

Senior AI/ML Engineer

FRAME PAYMENTS · REMOTE · OCT 2025 - JAN 2026

Sole AI/ML engineer shipping agentic document intelligence and multimodal compliance screening for merchant onboarding at a high-risk-payments platform.

- Deployed submission-time document-category validation using PydanticAI, Amazon Bedrock AgentCore, and Claude, surfaced as in-browser feedback for confidential uploads; reduced analyst follow-ups and resubmission requests ~90% by blocking miscategorized documents before they entered review queues.
- Built restricted-item compliance screening that combines product images with structured submission metadata; translated banking-partner policy requirements into automated checks supporting onboarding across tens of merchants, multi-location operations, and catalogs with hundreds of items per location.

Associate III, Senior Software Engineer — Data Integrations

JPMORGAN CHASE & CO. · NEW YORK, NY · MAY 2023 - APR 2025

Built enterprise data pipelines and analytics products for credit trading and risk teams while helping scale internal AI developer tooling.

- Pioneered firm-wide rollout of generative AI development tools (LLMs, GitHub Copilot, Codeium) as one of the first 500 engineers, delivering feedback that shaped secure, scalable deployment and leading internal workshops to boost productivity for technical and non-technical teams.
- Devised ML-driven anomaly detection framework for data health monitoring, reducing false positives by 25% and cutting manual review time, empowering risk teams to focus on critical alerts.
- Developed resilient Python wrapper for KDB+, democratizing data access for 5,000+ users with CI-style tests and optimized Q functions, reducing manual processing by over 10 hours per week.

Senior Software Engineer

STEALTH WEB3 STARTUP · REMOTE · MAR 2022 - JUL 2022

- Architected a scalable NFT-collateralized lending DApp using React, Next.js, TypeScript, MobX, and ethers.js, enabling secure, performant on-chain loans.
- Implemented end-to-end CI/CD with GitHub Actions, Husky, Cypress, and Jest to halve release cycles and raise test coverage to 90%, while mentoring a front-end intern to build key UI components.

Founder, Chief Technology Officer

SANDBLABS · REMOTE · JUN 2021 - MAR 2022

- Engineered and launched a Next.js/TailwindCSS marketing site with top-tier Lighthouse scores and secured \$75K in OCEAN tokens via a competitive DAO grant.
- Engineered Python pipelines to extract and normalize a 1M+ Web3-related Reddit posts & comments dataset and published the data on Ocean Protocol's Data Marketplace.

Demand Forecasting Extern

GAP INC. · SAN FRANCISCO, CA · JAN 2020 - MAY 2020

- Applied time-series decomposition and bias correction techniques to demand forecasts, boosting accuracy by up to 25% and minimizing surplus inventory costs.
- Owned the data engineering layer by architecting scalable MySQL and Apache Spark pipelines across 5+ years of sales data, improving data processing throughput by 95% and standardizing insights across forecasting teams.

PROJECTS

AI Agent Harness Configs

JAN 2026 - PRESENT · GITHUB

Cross-harness config set for Claude Code, GitHub Copilot, and related agent runtimes, packaging reusable instructions, tools, skills, and MCP presets.

- Built 40+ reusable instruction and skill packs plus shared harness configs that install across Claude Code, GitHub Copilot, Cursor, and related agent runtimes.
- Curated portable tool and MCP presets for code review, security, infrastructure, research, and workflow automation across environments.

MCP-Crawl4AI

APR 2025 - PRESENT · GITHUB

FastMCP server wrapping Crawl4AI so AI agents can crawl and extract web content through four focused tools with LLM-friendly markdown, cleaned HTML, or plain-text output.

- Built a FastMCP v3 server with lifespan-managed headless Chromium, session-aware crawling workflows, and automatic Playwright browser setup—only 2 runtime dependencies.
- Exposed 4 tools, 2 resources, and 3 prompt templates with both stdio and HTTP transport, enabling drop-in integration with any MCP-compatible AI agent.

NBA Basketball Database

APR 2021 - JUN 2021 · GITHUB

Open-sourced a comprehensive NBA analytics dataset covering 65K games and 13M play-by-play events, reaching 400K+ views and 45K+ Kaggle downloads.

- Implemented a high-performance, fully validated ETL pipeline using Python, Pandas, multiprocessing, pandera, pydantic, sqlalchemy, and SQL to extract data from the NBA's public API and update the dataset's SQLite database and CSV/JSON flat files with new data on a daily basis.

ProxyWhirl

MAY 2025 - PRESENT · GITHUB

Intelligent rotating proxy system with 9 rotation strategies, 114 auto-validated proxy sources, and self-healing capabilities that auto-eject dead proxies and implement circuit breakers.

- Implemented 9 rotation strategies (round-robin, weighted, geo-targeted, performance-based, and more) with configurable fallback chains and per-source health scoring.
- Built a self-healing pipeline that continuously validates 114 built-in proxy sources, auto-ejects unresponsive proxies via circuit breakers, and re-promotes recovered sources.

SKILLS

PROGRAMMING & SCRIPTING LANGUAGES: Python · SQL · TypeScript · JavaScript · Go · Rust · Bash

MACHINE LEARNING AND GEN AI: scikit-learn · Hugging Face Transformers · OpenAI API · LangChain · LangGraph · MCP (Model Context Protocol) · RAG · PydanticAI · Amazon Bedrock · Claude · Claude Code · GitHub Copilot · DeepAgents · AI Agent Skills Protocol

DATA ENGINEERING AND ANALYTICS: Pandas · NumPy · SQLAlchemy · Pydantic · Pandera · Apache Spark · Amazon Kinesis · AWS Glue · KDB+ · PostgreSQL · DynamoDB · ETL · SQLite

CLOUD INFRASTRUCTURE AND DEVOPS: AWS · Docker · Kubernetes · Terraform · GitHub Actions · Serverless Architecture · AWS Lambda · Google Cloud Functions · ECS · EKS · Prometheus · Grafana

EDUCATION

Bachelor of Science

University of California, Berkeley – College of Engineering, Industrial Engineering & Operations Research

AUG 2014 - DEC 2020

High School Diploma

The Hotchkiss School

AUG 2012 - MAY 2014

CERTIFICATIONS

AWS Certified Cloud Practitioner

Amazon Web Services (AWS) · Nov 2024

IBM Data Science Professional Certificate

IBM · Feb 2021