Kyle Johnson

Richmond, VA · gkjohns@gmail.com · kylejohnson.ai · linkedin.com/in/gkjohns · github.com/GKjohns
USMC veteran (Sergeant, E-5)

AI architect and strategist at Wells Fargo. Before that, two years founding Monument Labs, an AI studio shipping agentic products end to end (a live SaaS with 400+ users, an autonomous analyst agent in production, client platforms delivered), after six years at Meta (Senior Data Scientist, IC5) at billion-user scale. I run the full loop: discovery with non-technical stakeholders, architecture, evals, security posture, and the executive readout. Columbia CS & Statistics.

EXPERIENCE

Executive Director, AI Architect and Strategist · Wells Fargo

Sept 2026-Present · Richmond, VA · AI Product Acceleration & Innovation

- Individual contributor setting strategy, reference architecture, standards and working prototypes for agentic AI on the business side of Commercial Banking and Corporate & Investment Banking.

Founder · Monument Labs

July 2024-Present · Richmond, VA · AI R&D lab and consulting studio

- Built **The Analyst**, an autonomous AI data scientist running a production analytics function unattended: an agent loop on daily/weekly crons in an egress-locked sandbox, a six-tool MCP server (read-only SQL via a least-privilege Postgres role, ads, allowlisted git forensics), a durable findings ledger, and recommendations graded against measured outcomes. Its analyses drove activation from ~11% to 38%.
- Built and launched **Daylight Legal** (www.daylight.legal), a voice-first documentation SaaS: 400+ users, Stripe billing, paid acquisition. Whisper transcription → LLM structured extraction → court-ready records, 19 background workers (OCR, vision, report generation, lifecycle email), all LLM traffic through one model-picks chokepoint.
- Built **Margin** (marginfordata.com), a multi-tenant cloud data-science platform: a from-scratch sandboxed kernel-execution service (isolated subprocesses, resource limits, Docker) and an embedded agent that writes, runs, and debugs code in the user's live Python kernel, reading its own charts via vision.
- Delivered an AI operations platform for a nonprofit client: meeting-transcript extraction of decisions and action items, a tool-calling assistant over org data with row-level security, nightly cross-meeting insight synthesis, and a 15+-tool chief-of-staff agent in Discord.
- Served as interim **Head of Data Science** (consulting) for a ~$2B public company: assessed the function, delivered a transformation roadmap to the CEO, and coached the IC team toward strategic-partner positioning.

Senior Data Scientist (IC5) · Meta

2018-2024 · Seattle / Remote

- **Reality Labs (Trust):** built an LLM-based brand-trust measurement system classifying public posts into proactive-trust and implicit-distrust rates, reported weekly to the C-suite. Among the earliest LLM-labeling-as-metric work at Meta.
- **Counter-Terrorism:** designed an iterative network-propagation algorithm detecting coordinated militia groups that evaded text-based classifiers; accounted for ~40% of takedowns in the violent-militia category at low false-positive rates.

Sergeant (E-5) · United States Marine Corps

2009-2014 · Okinawa, Japan · Led teams and managed logistics operations, 2nd Assault Amphibian Battalion and Combat Logistics Regiment 37.

EDUCATION

B.A., Computer Science & Statistics · Columbia University · Minor in Mandarin Chinese

SKILLS

AI engineering & architecture: agentic systems, MCP servers, tool orchestration, prompt engineering, eval harnesses, agent security (sandboxing, egress control, least-privilege tooling), multi-provider routing, cost/latency budgeting · **Stack:** Python, SQL, TypeScript/Nuxt/React, Postgres/Supabase, Inngest, Vercel, Docker, Claude Agent SDK · **Statistics & ML:** causal inference, experiment design, Bayesian modeling, metric design