

Matt Coburn

AI executive who builds revenue-generating products in regulated, data-heavy enterprises

I've spent my career building AI products that ship in regulated industries and drive measurable revenue — from optimizing **\$200M+ in annual spend** at Expedia to shipping **KYC/AML compliance systems** as **VP of Data Science** at WorkFusion to generating **\$2M in enterprise revenue** as a founder serving MasterCard, Fannie Mae, and Freddie Mac.

As a VP, I've led teams of 15+ data scientists and engineers. As a founder, I built from zero to SOC 2-certified enterprise product with Fortune 100 clients. As a technologist, I've shipped production ML at 2 of the top 3 global automakers and built compliance screening products for financial services.

My domain runs deep in **compliance, document intelligence, and real estate**: I started my career originating mortgages, built AI products for mortgage-backed securities clients (Fannie Mae, Freddie Mac), and have spent the last decade building document extraction and screening systems for regulated enterprises. I believe AI will fundamentally transform how real estate operates — from leasing to screening to revenue management — and I want to lead that transformation.

Aristotle — Founding Engineer & Tech Lead

Los Angeles, CA | June 2024 - Present

Leading technical direction for an enterprise AI platform in production at **2 of the top 3 global automakers**. Building agentic AI systems that automate complex business workflows across regulated environments.

Key contributions:

- **Agentic AI architecture**: Built tool-using agents as explicit state machines with boundary validation, retry/escalation semantics, and human-in-the-loop controls — enabling enterprise adoption in regulated settings
- **Enterprise document intelligence**: Structured extraction pipeline with click-to-source provenance, schema validation, and auditable decision trails
- **AI governance in practice**: No model output reaches a user or changes system state unless it is schema-valid, auditable, and traceable to source documents
- **Team leadership**: Hired and led 4 engineers; established CI/CD, code review standards, release gates, and observability practices
- **Hybrid retrieval & evaluation**: BM25 + dense retrieval tuned on Precision@K / Recall@K with regression harnesses as release gates

Impact: Enabled enterprise workflows that reduced manual document review, enforced grounding guarantees, and unblocked adoption in regulated environments.

WorkFusion — VP of Data Science

New York, NY | 2023

Led **~15 data scientists** shipping compliance AI products in regulated financial environments. Owned the applied AI roadmap for KYC/AML, document extraction, and screening systems.

- **Compliance AI at scale**: Shipped transaction monitoring, risk scoring, and case prioritization systems for major financial institutions; improved detection outcomes by **~35%**
 - **AI screening & document intelligence**: Built LLM-based document extraction and auto-labeling for high-volume screening workflows with strict governance and explainability requirements
 - **Responsible AI operations**: Designed security-first MLOps, evaluation, and audit practices for production systems in regulated environments — bias mitigation, model monitoring, and compliance reporting
-

Tangible Intelligence – Founder & CEO

Dallas, TX | Jan 2020 – Apr 2023

Founded and led a document intelligence company serving Fortune 100 clients. Generated ~\$2M enterprise revenue with clients including MasterCard, Fannie Mae, and Freddie Mac.

- **Enterprise document extraction platform:** Built a no-code, HITL extraction system with ontology-driven schemas and click-to-source provenance – directly applicable to lease processing, tenant applications, and screening documents
 - **SafeScan – real-time compliance screening product:** Built a sensitive-data detection and anomaly screening system combining ML, rules, and search. Generated ~\$900K in revenue as a standalone product for enterprise compliance use cases
 - **Real estate & mortgage domain:** Served Fannie Mae and Freddie Mac, building AI systems for mortgage-backed securities document workflows – extraction, classification, and compliance validation
 - **Enterprise operations:** SOC 2-certified deployments, Fortune 100 enterprise sales, and an 8-person engineering and data science team
-

M Science (Jefferies subsidiary) – Principal Data Scientist, Founding Lead

New York, NY | May 2018 – Dec 2019

Built an internal AI startup at a major investment bank. Conceived, architected, and shipped an analytics platform to production for hedge fund clients including BlackRock, Two Sigma, and Citadel.

- **Enterprise analytics platform:** Built an Ontology-as-a-Service layer modeling companies, securities, events, and relationships across public and proprietary data
 - **Entity resolution & data quality:** Designed deterministic identity resolution and relationship semantics prioritizing reproducibility and auditability
 - **Client-facing AI product:** Built an application layer enabling institutional clients to apply shared analytics to private datasets with isolation and fine-grained access control
-

Expedia Group (Hotels.com) – Data Scientist

Dallas, TX | Oct 2016 – May 2018

Built large-scale revenue management and optimization systems influencing \$200M+ in annual advertising spend.

- **Revenue optimization at scale:** Built models estimating expected value of clicks across millions of keywords, enabling automated bid optimization across publisher inventory
 - **ROI & spend allocation:** Identified systemic underperformance across publisher inventory and reallocated spend, yielding ~240% ROI improvement on targeted segments
 - **Production pipelines:** Built revenue-critical production systems with monitoring, safeguards, and alerting
-

Earlier Career

Mortgage Industry | Dallas, TX | 2005 – 2008

Originated residential mortgages at Wells Fargo and an independent brokerage. This early career in real estate lending built deep domain understanding of mortgage workflows, underwriting, compliance requirements, and the document-heavy processes that define the industry – context that has directly informed my work building AI document intelligence and compliance systems.

Technology Stack

Core: Python, SQL, TypeScript, C, Zig **Backend & Infrastructure:** FastAPI, Pydantic, PostgreSQL, Redis, Docker, Kubernetes, AWS **Data & Analytics:** Pandas, NumPy, Spark, ETL/streaming pipelines, vector databases **AI/ML:** PyTorch, HuggingFace, OpenAI/Anthropic APIs, RAG architectures, hybrid retrieval, agentic frameworks

Specialties: Document intelligence, compliance AI, structured extraction, screening systems, AI governance, production ML, evaluation frameworks

Education

University of Texas at Dallas

- B.S. - Electrical Engineering, 2013
 - Masters Coursework - Computer Science, 2016
-

Building the Future of AI in Enterprise

I build AI products that generate revenue in regulated, data-heavy enterprises — and I've done it as a founder, as a VP, and at massive scale. I'm looking for VP-level AI leadership roles where I can drive product strategy, build high-performing teams, and deliver AI systems that transform how enterprises operate.