

Christopher Daigle

Harrisburg, PA | [linkedin.com/in/christopherdaigle](https://www.linkedin.com/in/christopherdaigle) | quantchris.com

AI & Data Science Executive | Insurance & Financial Services

Generative & Agentic AI • Natural Language Processing • Machine Learning • Statistical Modeling

Organizational Enablement • AI Governance & Evaluation • Executive Partnership

Twelve years taking quantitative systems into production, from econometric research and supervised modeling through natural language processing to multi-agent AI, six of them leading AI in P&C insurance. Sets the standards and evaluation methods that make AI delivery repeatable across an organization. Treats changing the process around a system as part of building it.

PROFESSIONAL EXPERIENCE

Arch Capital Services LLC

Sep 2024 – Present

Associate Vice President, AI Engineering

Apr 2026 – Present

Director, AI & Automation Engineering

Sep 2024 – Apr 2026

Direct five data scientists delivering AI products across six business units. Partner with seven functional leads spanning eight functions and roughly 150 staff in the US and UK, without line authority.

- **Submission Intake:** Reduced submission data entry from 30 minutes to about one minute across six business units, creating roughly 40 FTE of capacity redeployable to underwriting review and clearance, with 40% of Excess & Surplus Casualty submissions needing no human touch
- **Coding Agent Center for Enablement:** Founded and lead the C4E, setting agent standards, personas, and reusable components adopted across eight functions and closing engineering practice gaps for staff of varied technical depth
- **Development Acceleration:** Built an automated prompt engineering and evaluation capability that cut extraction schema development from four months to three hours, since adopted by five teams including claims data science and non-technical project management
- **Decision Frameworks:** Define the validation standards and confidence thresholds determining when an AI decision is automated, reviewed, or escalated, giving underwriting and operations a defensible basis for trusting automated output
- **Process Change:** Redefined enterprise clearance procedures with operations around automated intake, running both in parallel through cutover and shifting analyst roles and measures from data entry to exception review
- **Executive Partnership:** Represent Strategic Analytics to the C-suite across Operations, IT, Underwriting, and AI; scope initiatives for feasibility, resourcing, and ROI to shape multi-year roadmaps and guide platform and vendor investment

The Hartford Financial Services Group, Inc.

Jul 2020 – Sep 2024

Director of Data Science, Generative AI Factory (Organization Lead)

Dec 2023 – Sep 2024

Founding leader of the enterprise generative AI organization. Seven direct reports across data science, data engineering, product, prompt engineering, and MLOps. \$14M annual program budget; systems serving 22,000+ employees.

- **Productionization:** Scaled the Middle and Large Commercial underwriting guidelines RAG application from MVP to enterprise flagship product, establishing the repeatable pattern subsequently applied to other lines of business
- **Operating Model & Governance:** Established the federated and centralized GenAI operating model, science and technology standards, and the governance framework for ethical deployment; negotiated the roadmap across line-of-business leaders
- **Cost Performance:** Reduced planned annual capital spend from \$2.4M to roughly \$2,400 through model selection and retrieval design decisions while exceeding quality targets

Director of Data Science, Artificial Intelligence Factory

Mar 2023 – Dec 2023

Senior Data Scientist, Artificial Intelligence Factory (NLP)

May 2022 – Mar 2023

Technical lead for natural language processing and generative AI.

- **First Production GenAI:** Delivered the company's first generative AI system in production in four weeks, from no approvals to a working MVP rather than a proof of concept, clearing first-of-their-kind licensing, legal, and governance review for open-weight foundation models; selected over a parallel external consultancy build and appointed founding lead of the enterprise GenAI organization
- **Capability Building:** Led a team new to the domain, building foundation model, prompt, and cloud engineering capability as the work proceeded; every contributor advanced in role afterward
- **Applied NLP:** Deployed sentiment analysis for investor communications in one week and the Voice of the Customer topic model, increasing survey review coverage 10x and volume reviewed 50x
- **Teaching & Advisory:** Trained individual contributors through senior leaders and presented to enterprise audiences of 500+; evaluated external AI ventures for the strategic investments function

Senior Data Scientist, Claims Data Science (Workers Compensation)

Jul 2020 – May 2022

- **Model Portfolio:** Developed four and maintained nine supervised models predicting litigation likelihood, fraud, severity, disability duration, and complexity, informing reserving, legal spend forecasting, and claim routing; concentrated nurse

case management on the highest-complexity 5% of claims and identified likely fraudulent claims averaging roughly \$50,000 each

- **Practice Adoption:** Identified the absence of a reusable model development pattern and drove adoption of an MVP framework across data science teams, reducing initial model build from over one month to one day

Pratt & Whitney

Jan 2019 – Jul 2020

Manager, Data Scientist

May 2019 – Jul 2020

Analytics Developer

Jan 2019 – May 2019

- Established the analytics department for Commercial Engines Aftermarket Supply Chain

Boise Analytics

Dec 2017 – Dec 2018

Partner, COO & Chief Data Scientist

Founding leadership of the analytics practice

- Led 20 data scientists and owned 43 B2B client relationships, with accountability for delivery, hiring, and client outcomes

SELECTED PRODUCTION SYSTEMS

Submission Intake Platform | [Arch](#) • Multi-agent extraction, validation, and routing of broker submissions across six business units

Scale: roughly 200,000 submissions annually; extended to cyber loss runs, automobile and workers compensation schedules

Approach: multi-agent orchestration (LangGraph, Playwright), intelligent document processing (Cytora), Azure AI Foundry, and Databricks, with a purpose-built harness validating output consistency and correctness

Coding Agent Center for Enablement | [Arch](#) • Standards, agent personas, and reusable components for coding agent adoption across Strategic Analytics

Approach: harness-agnostic component design across GitHub Copilot, Claude Code, Codex CLI, Databricks Genie, and Snowflake Cortex

Underwriter Conversational Search | [The Hartford](#) • Retrieval augmented generation over thousands of Middle and Large Commercial guidelines containing tables, images, and color-encoded meaning

Outcome: guideline research effort reduced 95%, retrieval from roughly 10 minutes to seconds, seven-second latency, 95%+ positive user feedback

Approach: locally hosted FLAN-T5-XXL with quantization, migrated to GPT-3.5 with embeddings from ada-002 to text-embedding-3-large; ran cross-cloud because Azure held sole enterprise access to OpenAI models while the platform ran on AWS; cost reduced through context engineering, retrieval-tuned chunking, scoped prompts, and caching of known answers

Workers Compensation Triage Suite | [The Hartford](#) • Nine models scoring claims for litigation likelihood, fraud, severity, disability duration, and handling complexity

Approach: GLM (logistic, gamma, Tweedie), gradient boosting, random forest, clustering and anomaly detection, with ICD-10 codes engineered into predictive features

Aftermarket Supply Chain Analytics | [Pratt & Whitney](#) • Spend optimization and commodity classification for Commercial Engines procurement

Outcome: \$29M annual cost avoidance across 5,000+ vendors and 75,000+ components; classification coverage from 60% to 90% of \$16B in purchase orders at 94% F1, turnaround from eight hours to real time

Approach: constrained allocation modeling; NLP classification with TF-IDF and ensemble methods

RECOGNITION

TDAC Visionary Award, The Hartford (2023) – top 2% of employees, for the enterprise's first production GenAI capability

Innovation Award, Pratt & Whitney – five of approximately 240,000 UTC employees annually

Claims and Analytics Insight Boards, The Hartford – the Volatility Index, an algorithm surfacing overlooked high-cost claims, and deployed NLP work

SKILLS

Leadership: Enterprise AI Strategy • Organizational Enablement • Change Management • Use-Case Selection & Investment Evaluation • Vendor Evaluation • Talent Development

AI & Generative Systems: Agentic & Multi-Agent Systems • Foundation Models & Transfer Learning • RAG • Model Evaluation • MLOps

Statistical & Machine Learning: Econometrics • Statistical Inference • Time Series & Forecasting • GLMs • Tree Ensembles & Boosting • Clustering & Dimensionality Reduction • Experimental Design

Domain & Governance: P&C Insurance • Underwriting • Claims • Submission Lifecycle • Responsible AI Governance • Model Risk Controls

Platforms: Python • SQL • AWS • Azure AI Foundry • Databricks • LangGraph • Coding Agents • Hugging Face • PyTorch

LEADERSHIP FOUNDATION

Nine years in military and defense aviation, from sergeant flying large-scale unmanned aircraft in Iraq and Afghanistan to GS-15E (O-6 equivalent) lead pilot accountable for flight operations, personnel, and logistics on international deployments

EDUCATION

MS, Quantitative Economics (STEM) – University of Connecticut; two years PhD coursework on fellowship

BA, Economics, Quantitative Emphasis – Boise State University Honors College; minors in Math and Applied Math