

Idilio Moncivais, PhD

Westminster, CO | (303) 999-0001 | idilio@moncivais.com | linkedin.com/in/idiliomoncivais

PROFESSIONAL SUMMARY

Senior Data Science and Analytics Executive with 25+ years of hands-on research, applied ML/AI development, and enterprise platform leadership across Fortune-100 organizations. PhD in Applied Statistics — a working research scientist who personally designs and codes advanced ML models, time series and probabilistic forecasting systems, and AI agent architectures alongside executive strategy responsibilities. Founded and scaled Lumen's Data Science organization (trademarked 'FinancelQ'), governing core enterprise datasets spanning Salesforce, SAP, HRIS, SCM, Treasury, and FP&A; built and deployed production AI agents for enterprise customer service; and authored ten peer-reviewed research publications and three U.S. patents on AI systems, LLM deployment, data governance, probabilistic forecasting, and large-scale analytics. Leads a global, multi-region AI Data Engineering organization spanning the United States, Poland, Argentina, and India, navigating diverse regulatory and cultural environments while maintaining unified engineering standards. Combined CTO and CIO executive oversight as EVP at Unison Systems, unifying technology strategy and IT operations under a single accountable executive. Deep expertise across columnar and vector database ecosystems, Palantir Foundry and Ontology, and modern cloud-native data platforms — combining hands-on data engineering depth with executive platform strategy and C-suite partnership to deliver AI capabilities that create measurable competitive advantage.

CORE COMPETENCIES

Applied ML/AI Research & Development | LLM Integration & Agentic AI Systems | Time Series & Probabilistic Forecasting (ARIMA/SARIMA/Monte Carlo)

Global & Multi-Region Team Leadership (USA Poland Argentina India) | Combined CTO/CIO Executive Oversight | MLOps & Production Deployment

Columnar & Vector Database Architecture (Snowflake BigQuery Redshift Vertica) | Palantir Foundry & Ontology

Enterprise Data Platform Strategy | Data Governance, Quality & Compliance | Cloud Architecture (AWS Azure GCP Databricks)

AI Agent Design & Deployment | NLP & Conversational AI | Engineering Standards (Six-Sigma CMM ISO 8001 SDL)

PROFESSIONAL EXPERIENCE

Head of Financial Intelligence & Data Architecture | Lumen Technologies - Westminster, CO
2015 - Present

Founding leader of Lumen's Data Science organization — built from the ground up into a trademarked capability ('FinancelQ') recognized for groundbreaking research in classification analysis and AI-driven financial intelligence. A hands-on research scientist and executive who personally writes production Python, SQL, and R code; designs and deploys ML models; architects AI agent systems; and governs a modern data platform spanning columnar databases, vector database extensions, and Palantir Foundry — while owning enterprise data platform strategy, governance, and C-suite partnership for a global telecommunications organization.

- Built and lead a global, multi-region AI Data Engineering organization spanning the United States, Poland, Argentina, and India — managing distributed teams across four countries and three

continents, navigating diverse regulatory environments (U.S., EU/GDPR, LATAM, APAC) while maintaining unified engineering standards, delivery cadence, and data governance practices.

- Designed and operated enterprise columnar database environments over ten years — governing analytical workloads across Vertica, Google BigQuery, Amazon Redshift, and Snowflake, optimizing query performance, cost efficiency, and analytical throughput for high-volume financial and operational data at enterprise scale.
- Implemented and operationalized vector database capabilities over five years using native extensions from Google BigQuery, Snowflake, MongoDB, and PostgreSQL — enabling semantic search, similarity retrieval, and embedding-based analytical workflows that power AI agent knowledge retrieval and LLM-grounded analytics across the enterprise.
- Architected and operated Palantir Foundry and Ontology environments over two years — building semantic data models, ontology-driven data pipelines, and Foundry-native analytics applications that connect enterprise data assets into a unified, governed, and queryable knowledge layer for operational and executive decision-making.
- Architected and deployed enterprise AI customer service agents over four years — building 'Chat With Your Data' systems with backend integration connecting 4 quoting systems, 16 ordering systems, and 8 billing systems into a unified intelligence layer, enabling agents to surface actionable insights and create tasks autonomously in Salesforce, Gainsight, and other peripheral systems.
- Designed and deployed production LLM-based agent workflows using OpenAI, Anthropic (all model versions), Google, and Microsoft Copilot — evaluating commercial and open-source model ecosystems for enterprise suitability, conducting comparative benchmarking, and operationalizing models within governed, production-grade data environments.
- Personally designed and coded advanced ML models in Python, SQL, and R — including custom ARIMA/SARIMA time series forecasting systems, probabilistic forecasting using Monte Carlo simulation, categorical data analysis using likelihood statistical methods, and classification models for customer, product, and price cluster analysis.
- Founded and scaled Lumen's Data Science organization from inception — establishing research methodology, ML engineering standards, model development lifecycle practices, and an experimental culture that produced trademarked 'FinancelQ' classification research recognized across the organization for its analytical innovation.
- Fine-tuned and optimized ML models for dedicated enterprise tasks — designing evaluation frameworks, experiment methodologies, and model comparison workflows to ensure production models delivered measurable improvements in forecast accuracy, classification precision, and decision quality.
- Defined multi-cloud data platform strategy on Databricks, Azure, GCP, and AWS — architecting data lakehouse environments with streaming and batch pipelines, data cataloging, observability standards, and quality frameworks enabling both analytics and AI agent use cases at enterprise scale.
- Established enterprise data governance and AI governance frameworks — defining data stewardship accountability, model risk management standards, responsible AI practices, and compliance infrastructure under GDPR, CCPA, HIPAA, and CPNI in regulated production environments.
- Directed SAP S/4 HANA ERP modernization — leading migration and integration with analytics environments, strengthening the integrity and consistency of financial intelligence data across the enterprise.
- Own data architecture and governance across core enterprise datasets spanning Salesforce, SAP, HRIS, SCM, Treasury, and FP&A — providing trusted, governed data foundations for financial intelligence, workforce analytics, supply chain visibility, and enterprise planning across the FinTech organization.
- Serve as part of Lumen's finance leadership, responsible for reporting on EBITDA and ARPU for company-wide visibility into financial goals compliance — delivering trusted, governed financial metrics that inform executive and Board-level decision-making.

- Operated under rigorous engineering standards — Six-Sigma, CMM, ISO 8001, and SDL frameworks governing data science and software development practices across the organization.
- Secured two U.S. patents for novel data integration and large-scale analytics inventions, and published ten peer-reviewed research papers on enterprise AI, LLM deployment, agentic analytics, Databricks platform optimization, Unity Catalog governance, probabilistic forecasting, applied statistics and biostatistics methodology, and hallucination prevention in production AI systems.
- Partnered directly with C-suite leaders as strategic thought partner — translating complex ML research, AI agent capabilities, and platform strategy into clear executive narratives that shaped enterprise investment decisions.

Adjunct Professor, Applied Statistics & Data Science | University of Northern Colorado - Greeley, CO 2020 - 2025

Graduate faculty appointment in applied statistics and data science — maintaining active research engagement and mentoring emerging data scientists alongside executive practice.

- Taught graduate-level Data Science, Applied Statistics, and Research Methods — instructing advanced statistical modeling, ML techniques, and research methodology to practitioners entering enterprise data science roles.
- Mentored students in Python, R, SQL, SAS, and SPSS with hands-on emphasis on statistical rigor, reproducible analysis, and applied ML model development for real-world business problems.

Executive Vice President & Chief Technology Officer | Unison Systems, Inc. - Greenwood Village, CO 2010 - 2015

P&L-accountable technology executive holding combined CTO and CIO executive oversight — leading applied data science, analytics platform delivery, enterprise IT operations, and engineering team management for a national consulting organization. Personally involved in technical architecture and ML development alongside executive leadership responsibilities.

- Held direct executive oversight of the CIO function and enterprise IT infrastructure alongside CTO responsibilities — unifying technology strategy, IT operations, and data/analytics leadership under a single accountable executive, a combined-scope mandate spanning the full breadth of enterprise technology and information systems.
- Led development of a high-volume Spectrum Analytics predictive platform — personally contributing to ML architecture design for automated customer experience scoring and maintenance prediction at scale, recognized with a granted U.S. patent.
- Defined and executed enterprise data science strategies for national clients — including ML model development roadmaps, data governance frameworks, and analytics engineering standards.
- Managed technical resource allocation for 300+ employees and contractors including nine direct reports, maintaining engineering quality standards across concurrent data science, analytics, and IT delivery programs.
- Established a Master Data Management practice across a 35-person matrixed team — standardizing data definitions, lineage documentation, and quality controls enabling trustworthy analytics and ML workloads.
- Designed Loyalty Program analytics for a large retail client — implementing behavioral segmentation, customer lifetime value modeling, and retention prediction using centralized transaction data.
- Governed Vertica-based enterprise data warehouse environments for large telecom clients — managing high-throughput analytical query workloads and establishing columnar database performance and optimization standards.

Sr. Director, Data Engineering, Management & Integration | Comcast Corporation - Littleton, CO 1993 - 2010

Senior data engineering and analytics leader governing enterprise-scale data infrastructure for a national telecommunications provider — processing 15 billion records nightly with strict engineering standards including SDL software development lifecycle practices established from this period forward.

- Oversaw enterprise data warehousing and BI environments processing 15 billion records nightly — governing data quality, pipeline reliability, and analytical performance under strict SDL engineering standards.
- Led a 200+ person Data Engineering organization, governing ETL, data integration, and analytical pipeline operations using Ab Initio, Informatica, Oracle, and Vertica.
- Directed engineering of subscriber identity management platform — an early-generation classification and authentication system securing millions of user identities and enabling customer behavioral analytics.
- Created nationwide GIS infrastructure integrating Oracle Spatial capabilities — enabling data-driven network planning and operational analytics at geographic scale.

EDUCATION

Doctor of Philosophy (PhD), Applied Statistics and Research Methods | University of Northern Colorado, Greeley, CO

Master of Science (MSc), Applied Statistics and Research Methods | University of Northern Colorado, Greeley, CO

Bachelor of Science, Military Administration (Physics & Mathematics) | Heroico Colegio Militar, U.D.E.F.A., Mexico City

CITI Program Certification — Research Ethics (Responsible Conduct of Research)

PATENTS & PUBLICATIONS

- US Patent No. 94/974,51B2: Cable Network Data Analytics System (CNDAS)
- US Patent No. 63/378,026: Enhanced Intake of BAI Files
- US Patent No. 63/380,887: Processing of Large Data Volumes (PLDV) from RDBMS
- Moncivais, I. (2026). Comparing Enterprise Data Analytics Capabilities of Databricks Genie versus Microsoft Copilot Agent: A Comprehensive Literature Review. Zenodo ©Lumen Technologies. doi.org/10.5281/zenodo.19687528
- Moncivais, I. & Fuentes, B. (2026). Guidelines to Preventing Artificial Intelligence Hallucinations in Microsoft Fabric. Zenodo ©Lumen Technologies. doi.org/10.5281/zenodo.19687703
- Moncivais, I. (2026). Towards a Hybrid Query Architecture in Lakehouse Environments: When to Use Databricks SQL Materialized Views vs. Interactive Spark Notebooks for Complex Analytical Workloads. Zenodo ©Lumen Technologies. doi.org/10.5281/zenodo.19701778
- Moncivais, I. (2026). From Insight to Action; Unlocking Real Value from Databricks Investments. Zenodo ©Statistica LLC. doi.org/10.5281/zenodo.20036490
- Moncivais, I. (2026). Guidelines for End-User Prompt Writing for Data Retrieval Using Databricks Genie. Zenodo ©Lumen Technologies. doi.org/10.5281/zenodo.20096925
- Moncivais, I. (2026). Optimizing Databricks Compute by Eliminating Waste, Reducing Costs, and Maximizing Performance. Zenodo ©Statistica LLC. doi.org/10.5281/zenodo.20284412
- Moncivais, I. & Moncivais, H. (2026). Scientific Forecast of the 2026 FIFA World Cup. Zenodo ©Statistica LLC. doi.org/10.5281/zenodo.20587172
- Moncivais, I. (2026). Databricks Unity Catalog: A Unified Governance Framework for Enterprise Data and AI Platforms. Zenodo ©Statistica LLC. doi.org/10.5281/zenodo.20603660
- Moncivais, I. (2026). Data Architecture and Governance in the Agentic Era: Paradigm Shifts from the 2026 Databricks Data + AI Summit. SSRN ©Statistica LLC. doi.org/10.2139/ssrn.6979203

- Moncivais, I. (2026). Two Statistical Cultures: How Domain, Stakes, and Regulation Shape Applied Statistics and Biostatistics. Zenodo ©Statistica LLC. doi.org/10.5281/zenodo.21398776

TECHNICAL PROFICIENCIES

Languages & ML: Python (advanced), R (custom analytical functions), SQL (complex), MATLAB, SAS, SPSS | ARIMA/SARIMA, Monte Carlo simulation, likelihood methods, classification

AI & LLMs: OpenAI, Anthropic, Google Gemini, Microsoft Copilot | LLM fine-tuning, prompt engineering, RAG, agentic systems, model evaluation

Columnar Databases: Snowflake, Google BigQuery, Amazon Redshift, Vertica | 10 years enterprise analytical workload optimization and governance

Vector Databases: Google BigQuery (vector search), Snowflake (Cortex vector), MongoDB Atlas Vector Search, PostgreSQL (pgvector) | 5 years embedding pipelines and semantic retrieval

Palantir: Foundry, Ontology, AIP | 2 years semantic data modeling, ontology-driven pipelines, and Foundry-native analytics applications

Data Platforms: Databricks, AWS, Azure, Google Cloud, Kafka, Airflow | Data lakehouse architecture, streaming pipelines, data cataloging

BI & Visualization: Tableau, Power BI, Microsoft Fabric, SAP Analytics Cloud, SAP BW | Self-serve analytics

Enterprise Platforms: SAP S/4 HANA, SAP BPC, HRIS, SCM, Treasury, FP&A, Anaplan, Hyperion, Celonis | Salesforce, Gainsight | Informatica, Ab Initio

Engineering Standards: Six-Sigma, CMM, ISO 8001, SDL | MLOps, model monitoring, CI/CD | GDPR, CCPA, HIPAA, CPNI, FERPA

Other: JavaScript, React | Agile delivery | Vendor and budget management | dbt-compatible transformation frameworks