

SUMMARY

AI Pioneer and Data Executive – People & Technology Leader – Data Scientist. Technically minded and highly analytical executive professional with extensive experience in applied research, true Data Science, P&L results, team building, data engineering, and advanced analytics. Stellar track record of success in artificial intelligence and machine learning, data migration, integration, architecture, and business intelligence. Proven background in developing and executing modern engineering practices of cloud-based data and analytics platforms. Highly skilled in forging and fostering trusted relationships at all levels. Well-versed in devising and executing data-driven analytical solutions. A solid history of success in identifying, implementing, and spearheading strategic initiatives and leading projects from inception to completion for the smooth running of operations. Leader of large global multi-disciplinary teams. Team builder.

DATA SCIENCE / DATA MANAGEMENT

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 forecasting systems, and AI agent architectures alongside executive strategy responsibilities. Founded and scaled Lumen's Data Science organization (trademarked 'FinanceIQ'), built and deployed production AI agents for enterprise customer service, and authored seven peer-reviewed research publications and three U.S. patents on AI systems, LLM deployment, data governance, and large-scale analytics. 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.

EDUCATION

The University of Northern Colorado
Doctor of Philosophy in Applied Statistics and Research Methods

Greeley, CO

The University of Northern Colorado
Master of Science in Applied Statistics and Research Methods

Greeley, CO

Heroico Colegio Militar, U.D.E.F.A.
Bachelor in Military Administration

Tlalpan, CDMX, Mexico

INDUSTRY EXPERIENCE

Head of Financial Intelligence and Data Architecture
Lumen technologies (formerly CenturyLink & Level3)

Mar 2015 to present
Denver, CO

Founding leader of Lumen's Data Science organization — built from the ground up into a trademarked capability ('FinanceIQ') 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.

- 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, categorical data analysis using likelihood statistical methods, and classification models for customer, product, and price cluster analysis, deploying these directly into production analytical workflows.
- 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 'FinanceIQ' 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.
- 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 seven peer-reviewed research papers on enterprise AI, LLM deployment, agentic analytics, Databricks platform optimization, Unity Catalog governance, 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 and Data Science

Jan 2020 to Present

University of Northern Colorado

Greeley, CO

- **Taught graduate-level courses** in Data Science, Applied Statistics, and Research Methods, equipping students with rigorous quantitative skills and practical tools for enterprise analytics and decision support.
- **Mentored emerging data professionals** on SQL, R, Python, SAS, and SPSS, emphasizing reproducible analysis, data governance awareness, and real-world problem framing.
- **Bridged academic and enterprise practice** by teaching quantitative, qualitative, and mixed-methods research, helping students translate complex findings into actionable insights for organizations.

Executive Vice President and CTO
Unison Systems, Inc.

Apr 2010 to Mar 2015
Greenwood Village, CO

P&L-accountable technology executive leading applied data science, analytics platform delivery, and engineering team management for a national consulting organization. Personally involved in technical architecture and ML development alongside executive leadership responsibilities.

- 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 and analytics 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.

Senior Director of Data Engineering, Management and Integration
Comcast Corporation

Aug 1993 to Apr 2010
Englewood, CO

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.

ACADEMIC EXPERIENCE

Adjunct Professor of Advanced Analytics
The University of Northern Colorado

Aug 2020 to present
Greeley, CO

Teaching graduate school courses in data science, applied statistics, and research methods, including SQL, R, SPSS, SAS, MathLab, Python, and NVIVO, plus Quantitative, Qualitative, and Mixed research methods, in the following courses:

SRM523	Introduction to SQL	Spring 2023
SRM600	Introduction to Graduate Research	Summer 2023
SRM522	Introduction to R	Fall 2023
SRM523	Introduction to SQL	Fall 2023
SRM600	Introduction to Graduate Research	Fall 2023
SRM622	Direct Study: Advanced Data Analysis	Fall 2023
SRM600	Introduction to Graduate Research	Summer 2024
SRM600	Introduction to Graduate Research	Fall 2024
SRM600	Introduction to Graduate Research	Spring 2025
SRM523	Introduction to SQL	Spring 2025
SRM600	Introduction to Graduate Research	Spring 2025
SRM600	Introduction to Graduate Research	Summer 2025
SRM723	Advanced Statistical Consulting	Fall 2025
SRM600	Introduction to Graduate Research	Spring 2026

Teaching Assistant/Graduate Assistant
The University of Northern Colorado

Jan/2020 – Dec/2022
Greeley, CO

Assisting in teaching graduate school courses in data science, applied statistics, and research methods, including SQL, R, SPSS, SAS, MathLab, Python, and NVIVO, plus Quantitative, Qualitative, and Mixed research methods in the following courses:

SRM523	Introduction to SQL	Spring 2020	Dr. Larkins
SRM523	Introduction to SQL	Spring 2021	Dr. Larkins
SRM686	Qual Case Study Res	Fall 2021	Dr. Larkins
SRM700	Advanced Research Methods	Fall 2021	Dr. Larkins
SRM700	Advanced Research Methods	Summer 2022	Dr. Larkins
SRM700	Advanced Research Methods	Fall 2022	Dr. Larkins

PUBLICATIONS

Two Statistical Cultures: How Domain, Stakes, and Regulation Shape Applied Statistics and Biostatistics

Idilio Moncivais, Ph.D.

<https://doi.org/10.5281/> (TBD)

Jun/2026

This paper reframes the contrast between industry applied statistics and biostatistics as a difference in warrant, stakes, and regulation rather than a difference in mathematics. Extending Breiman's 2001 two-cultures framing, it argues that the methodological repertoires, software stacks, and pre-specification norms that distinguish the two communities are explicable from where each sit on the axes of evidence threshold, decision consequence, and regulatory codification. The analysis span philosophy, methodology, regulation, and tooling across both cultures; grounds itself in two cross-read case studies (customer churn prediction and oncology drug efficacy); and engages current frameworks from ICH-GCP through the NIST AI Risk Management Framework, ISO/IEC 42001, the EU AI Act, and Federal Reserve SR 11-7. Tracing an expanding set of convergence sites — real-world evidence, federated learning, machine learning as a regulated medical device, the data-leakage reproducibility crisis, the reckoning over statistical versus clinical significance, and AI risk governance — the paper argues that the boundary between the cultures is thickening rather than eroding: the two are building shared infrastructure rather than merging into one. It closes with reciprocal recommendations for what each culture has to learn from the other.

Databricks Unity Catalog: A Unified Governance Framework for Enterprise Data and AI Platforms

Idilio Moncivais, Ph.D.

<https://doi.org/10.5281/zenodo.20493399>

Jun/2026

This paper provides a comprehensive, vendor-independent systems analysis of Databricks Unity Catalog as an implementation of active, in-engine data governance. It examines how the catalog addresses structural tensions within distributed lakehouse architectures by embedding policy evaluation, deterministic lineage capture, and real-time credential vending directly inside the query runtime engine rather than relying on external, passive tools. Organized around six technical evaluation dimensions—governance effectiveness, scalability, security, administrative complexity, interoperability, and performance—the paper contextualizes Unity Catalog alongside traditional data catalogs (such as AWS Glue and Apache Atlas), peer commercial offerings (Snowflake Horizon), and emerging open-source standards like Apache Polaris and Apache Gravitino. Finally, it outlines industry adoption realities, exploring how its capabilities align with prominent global regulatory frameworks like the EU AI Act, GDPR, and the NIST AI Risk Management Framework, while providing strategic recommendations for multi-cloud enterprise platform architects.

Optimizing Databricks Compute by Eliminating Waste, Reducing Costs, and Max Performance

Idilio Moncivais, Ph.D.

<https://doi.org/10.5281/zenodo.20284412>

May/2026

The Databricks Efficient Compute Utilization (2026) framework is a comprehensive governance and engineering standard designed to eliminate wasteful cloud compute spending on Databricks. By replacing ad hoc infrastructure habits with automated, policy-enforced discipline, the framework spans a mandatory compute selection decision tree, Terraform-deployed cluster policies, a cost-optimized Medallion data architecture, and a structured FinOps review cadence. Supported by a 55-citation academic literature review, this research demonstrates that organizations implementing comparable frameworks consistently achieve 25% to 40% reductions in cloud compute spend while simultaneously improving pipeline performance, governance consistency, and cost transparency.

From Insight to Action; Unlocking Real Value from Databricks Investments

Idilio Moncivais, Ph.D.

<https://doi.org/10.5281/zenodo.20036490>

Apr/2026

his industry white paper introduces the Decision Insight-to-Action Framework (DIAF), a socio-technical model for closing the last-mile gap between analytics capability and business outcomes in modern lakehouse environments. Synthesizing recent peer-reviewed research across information systems, decision science, and organizational behavior, it identifies five recurring failure factors — trust, ownership, latency, accessibility, and cognitive load — that determine whether an insight becomes a decision, and traces the five-stage path (Insight → Trust → Ownership → Action → Feedback) every insight must traverse to produce measurable value. DIAF is operationalized through three executive-ready artifacts — a failure-point overlay, a one-page diagnostic heatmap, and a four-layer socio-technical stack view — and a 90-day roadmap covering diagnosis, remediation, and institutionalization. Written for senior data and business leaders responsible for translating Databricks investment into enterprise outcomes, the paper argues that the next wave of value will not come from faster queries or richer catalogs, but from making insights act on the world rather than merely describe it.

Guidelines for End-User Prompt Writing for Data Retrieval using Databricks Genie.

Idilio Moncivais, Ph.D.

<https://arxiv.org>

Mar/2026

This paper, develops a practical framework to help non-technical users effectively retrieve and analyze enterprise data through natural-language interfaces. The study explains how modern Natural Language Interfaces for Databases (NLIDBs), powered by large language models and tools such as Databricks Genie, allow users to query complex data systems conversationally but remain highly sensitive to how prompts are written. The paper synthesizes insights from prompt engineering, text-to-SQL research, and human-computer interaction to propose clear guidelines that improve prompt clarity, context specification, schema awareness, and analytical intent. It also outlines advanced prompting strategies—including structured prompts, iterative refinement, and example-based prompting—to help users handle more complex analytical tasks such as multi-table joins, aggregations, temporal analysis, and business-logic constraints. By translating technical research into accessible best practices, the paper argues that prompt-writing literacy is becoming a key skill for organizations seeking to expand reliable self-service analytics, reduce dependence on specialized data teams, and improve trust in AI-mediated data

access.

Towards a Hybrid Query Architecture in Lakehouse Environments. When to Use Databricks SQL Materialized Views vs. Interactive Spark Notebooks for Complex Analytical Workloads
Idilio Moncivais, Ph.D. <https://doi.org/10.5281/zenodo.19701778> Feb/2026

This technical position paper by Idilio Moncivais, Ph.D., published by Lumen Technologies in February 2026, argues that neither Databricks SQL materialized views nor interactive Spark notebooks represent a universally superior approach for complex analytical workloads in lakehouse environments. Instead, the paper advocates for a principled hybrid architecture comprising three layers — a declarative materialized view core, an imperative notebook exploration layer, and an adaptive bridge via Delta Live Tables — in which workloads are explicitly routed based on five criteria: query shape stability, data freshness requirements, workload complexity, governance obligations, and cost priority. The paper synthesizes peer-reviewed literature to substantiate key performance claims, including up to 25× query speedups via Adaptive Query Execution and a 2–4× DBU cost premium for notebook compute over SQL warehouse compute, while also extending the framework to address federated governance challenges across multi-lakehouse enterprise environments, regulatory compliance requirements under GDPR, SOX, and HIPAA, and emerging ESG disclosure obligations tied to cloud compute energy consumption.

Guidelines to Preventing Artificial Intelligence Hallucinations in Microsoft Fabric.
Idilio Moncivais, Ph.D. <https://doi.org/10.5281/zenodo.19687703> Jan/2026

This technical guideline by Idilio Moncivais, Ph.D., and Bernardo Fuentes, M.Sc., published by Lumen Technologies in February 2026, presents a comprehensive, multi-layered framework for mitigating AI hallucinations — fluent but factually incorrect or ungrounded outputs — in Microsoft Fabric's Copilot environment. The paper begins by categorizing hallucinations into two types (intrinsic, where outputs deviate from input context, and extrinsic, where outputs contradict real-world facts) and tracing their root causes to structural LLM limitations such as noisy training data, next-token prediction objectives, ambiguous prompts, and incomplete enterprise datasets. Against this backdrop, the authors propose a seven-layer defense-in-depth strategy encompassing governance and access controls, data quality management, retrieval-augmented generation (RAG) for grounding, structured prompt engineering using techniques like the ICE method and chain-of-thought prompting, human-in-the-loop validation, automated detection via Azure AI Content Safety, and continuous monitoring. The paper concludes with a phased implementation roadmap — starting with data quality and governance foundations before scaling toward advanced RAG pipelines and automated correction systems — and aligns the overall approach with Microsoft's Responsible AI principles, while acknowledging that hallucination elimination is not fully achievable and that ongoing human oversight remains a permanent organizational requirement.

Comparing Enterprise Data Analytics Capabilities of Databricks Genie versus Microsoft

Copilot Agent: A Comprehensive Literature Review.

Idilio Moncivais, Ph.D.

<https://doi.org/10.5281/zenodo.19687528>

Jan/2026

This paper analyzes how two emerging AI platforms are reshaping enterprise intelligence and decision-making. Drawing on more than 80 peer-reviewed studies, the paper compares the architecture, capabilities, performance, and governance requirements of Databricks Genie and Microsoft Copilot. The review explains that Genie represents the “query-less analytics” paradigm, translating natural-language questions into SQL to democratize access to data for non-technical users, while Copilot Agent embodies the rise of agentic AI systems capable of orchestrating multi-step workflows and autonomously executing operational tasks across enterprise domains. Through a detailed examination of system design, real-world performance metrics, governance challenges, and deployment strategies, the paper argues that the two platforms are not direct competitors but complementary technologies: Genie excels at enabling conversational data analysis, whereas Copilot Agent automates complex organizational processes. The study concludes with a decision framework and phased implementation roadmap, emphasizing that successful enterprise adoption depends less on the underlying AI technology and more on organizational readiness, strong data governance, and effective change management.

SAP Accounts Payable Integration using Imperative Execution in an Interactive Spark Notebook. Technical Reference Documentation

Idilio Moncivais, Ph.D.

<https://intranet.lumen.com>

Dec/2025

The FLH_INTEGRATION_SAP_AP pipeline is a Databricks-native data integration solution that automates the daily extraction, transformation, and semantic enrichment of nine SAP Accounts Payable tables — spanning document headers (BKPF), line items (BSEG), vendor master data (LFA1), payment terms (T052), open items (BSIK), cleared items (BSAK), payment run headers and line items (REGUH/REGUP), and payment instruments (PAYR) — from Lumen Technologies' legacy SAP R/3 ERP system into governed Delta Lake tables within Unity Catalog. Running each morning in under ten minutes on a mid-size SQL Warehouse, the pipeline applies consistent AP-specific scope filters (client 700, fiscal periods from January 2024 forward, posting keys 21/22/27/31/32) and annotates all 693 columns with business-readable metadata comments, creating the semantic bridge that enables Databricks Genie and Microsoft Copilot agents to translate natural-language questions into accurate SQL across the full accounts payable lifecycle — from initial invoice ingestion through vendor lookup, open-item aging, clearing, payment program execution, and final check or EFT materialization — without requiring end users in Finance, AP Operations, or Audit to write any code.

Databricks Unity Catalog: Thematic and Source Data Comprehensive Analysis. Technical Reference Documentation.

Idilio Moncivais, Ph.D.

<https://intranet.lumen.com>

Nov/2025

This document is a confidential Technical Reference Manual, presenting the design, execution,

and findings of an automated thematic and source-system analysis of Lumen's Databricks Unity Catalog estate conducted in March 2026. Using a five-stage Python pipeline (`uc_thematic_analysis_v8`), the analysis cataloged 6,138 data objects across eight Unity Catalog namespaces — spanning the Analytics Serving Layer, Enterprise Reference Layer, Data Serving Layer, Finance and FP&A Unified Data Layers, and a Finance landing zone — enriching each object with 90-day audit access statistics, column-level lineage, semantic embedding clusters (k-means, k=16), and LLM-generated domain labels via Claude Opus. The document's principal findings highlight an extreme concentration of read activity in a small set of high-value FP&A and production serving objects (65% of 6.4 million total reads concentrated in just three clusters), identify 17 upstream source systems spanning OSS, BSS, CRM, and ERP categories including Sonus, Kenan, Salesforce, SAP R/3/ECC, and SAP S/4HANA, and deliver governance recommendations covering LLM label re-runs, VIEW rationalization in the Data Serving Layer, catalog deprecation for fully dark namespaces, and prioritization of the FP&A hot-object core as the estate's primary AI-readiness and Databricks Genie configuration target.

Comprehensive Analysis of Lumen's 2026 Databricks Unity Catalog. Technical Reference Documentation

Idilio Moncivais, Ph.D.

<https://intranet.lumen.com>

Oct/2025

This document is a comprehensive technical reference for Lumen Technologies' Databricks Unity Catalog platform, analyzing four dimensions of the enterprise lakehouse deployment as of March 2026: physical storage architecture (11 external locations across a three-zone Landing → ASL → UDL data lake spanning two financial domains on Azure East US 2), data asset inventory (23,231 registered objects including 10,857 managed Delta tables, 9,628 Lakehouse Federation foreign tables from Oracle and SQL Server, 1,088 materialized views, and early-stage AI capabilities via four Vector Search indexes and Anthropic model serving), governance and security posture (102,401 privilege grants across 112 grantees with a prioritized risk register identifying critical issues such as wildcard grants to all workspace users on financial catalogs and unidentified service principals holding full privileges), and operational consumption (53,366 DBUs across 25 workspaces over a six-day billing window confirming an automation-first platform profile at 61.5% Jobs compute, with active Predictive Optimization, Lakeflow Connect ingestion, and a nascent RAG application stack). The document serves as both a platform audit artifact and a governance remediation roadmap, with prioritized recommendations addressing credential segmentation, access control tightening, observability gaps in audit logging and lineage, and cost attribution instrumentation.

Full Integration of Billing with Orders using Imperative Execution in an Interactive Spark Notebook. Technical Reference Documentation

Idilio Moncivais, Ph.D.

<https://intranet.lumen.com>

Sep/2025

The FLH_FULL_INTEGRATION Technical Reference documents a production Databricks SQL pipeline that daily constructs a single, governed Delta Lake table unifying Lumen Technologies' financial billing transactions — spanning 17 billing systems and 21 ordering

systems — with customer, product, and service order dimensions. The pipeline ingests data from both a Legacy ERP (January 2024 through September 2025) and a Modern ERP (October 2025 onward), applies seven sequential processing steps including dual-ERP unification via UNION ALL, three-layer deduplication to prevent cardinality fan-out, deterministic CUI data masking for controlled unclassified information accounts, and a single-pass LEFT JOIN with rounding-matched reconciliation, producing a 61-column AI-queryable dataset written to Unity Catalog at udl_dlh_finance.fin_integration.FLH_FULL_INTEGRATION. The v3.0 release resolved six critical defects — most notably a billions-dollar USD inflation caused by an unguarded cross-reference join fan-out and a non-deterministic masking function incompatible with GROUP BY deduplication — and the resulting table, enriched with business-friendly column metadata, serves as the authoritative data foundation for natural-language financial queries via Databricks Genie and Microsoft Copilot.

Enhanced Processing of Large Data Volumes from Inside Relational Databases

Idilio Moncivais-Pinedo Patent Number: US-63380887 Published Apr 25, 2024
<https://patentimages.storage.googleapis.com/be/e7/9e/39e8a91c1f1653/US20240134889A1.pdf>

This disclosure describes systems, methods, and devices for analyzing data stored in a relational database. A method may include installing a structured query language (SQL) server on a host server; installing statistical analysis modules on the host server; executing the statistical analysis modules within a relational database of the SQL server to analyze data stored in the relational database; and generating outputs based on the execution of the statistical analysis modules within the relational database.

Enhanced Intake of Bank BAI Files

Idilio Moncivais-Pinedo Patent Number: US-20240112278 Published Apr 4, 2024
<https://patentimages.storage.googleapis.com/e8/e2/ee/aa44192380787d/US20240112278A1.pdf>

This disclosure describes systems, methods, and devices for processing Bank Administration Institute (BAI) files. A method may include identifying, by at least one processor of an integration server, copies of BAI files indicative of payment information; importing, by at least one processor, data from the BAI files into a staging table; parsing, by the at least one processor, the data in the staging table using a hierarchy defining a file level, a group level, an account level, and a transaction level to generate hierarchized data; normalizing, by the at least one processor, the hierarchized data into tables; and copying, by the at least one processor, differences between the hierarchized data and previous hierarchized data to a workflow associated with identifying payment amounts in the hierarchized data.

Law Enforcement Efficacy in Domestic Violence Calls Based on the Responding Officer's Experience, Training, and Scholarly Level

Moncivais Pinedo, I. 2022. Available from Dissertations & Theses @ University of

Northern Colorado; ProQuest One Academic. (2771893682).

<https://www.proquest.com/docview/2771893682?pq-origsite=gscholar&fromopenview=true&sourcetype=Dissertations%20&%20Theses>

The purpose of this study was to provide insights into whether law enforcement officers (LEO) have been able to produce positive outcomes based on their responses when the victims have called 911. Law enforcement officers in this sample were grouped into three clusters based on the number of years of patrol experience (10 years or less versus more than ten years), the education level (less than a bachelor's degree versus at least a bachelor's degree), and training (advanced training versus basic training). The key findings were as follows: A) The three clusters of law enforcement officers were well distinguished. The first cluster consisted of law enforcement officers with ten years or less of patrol experience, who held a degree of less than a bachelor's, and who had received advanced training. The second cluster comprised law enforcement officers with more than ten years of patrol experience, who held at least a bachelor's degree and had received advanced training. The third cluster was composed of law enforcement officers with ten years or less of patrol experience, who had less than a bachelor's degree, and who only had basic training; B) The percentage of arrests given to the predominant aggressor party/parties by law enforcement officers with basic training only was significantly less than that of law enforcement officers with advanced training; C) The odds of giving an arrest to the predominant aggressor party/parties mattered less about their number of years of patrol experience and their education level but more about their certification levels. The result of these findings led to the following suggestions for law enforcement: Place a strong emphasis on encouraging officers to complete the advanced training, which could be done using any of the following strategies: (a) Conducting an educational campaign to highlight the importance of the advanced training certificates and motivate more officers to go through the advanced training, (b) Having important elements of the advanced training modules placed into the basic training modules to reduce the effect of failing to take advanced training modules, (c). Removing or reducing any existing structural barriers to taking the advanced training certificates that could prevent more officers from having them, and/or (d) making some modules of the advanced training required in addition to the basic modules.

Pandemic Poetry

Lahman, M. K. E., De Oliveira, B., Fan, X., Hodges, J., Moncivais, I., & Royse. 2021. *Qualitative Inquiry* 27(6), 752-756. <https://doi.org/10.1177/1077800420948099>

Members of a graduate educational ethnography course created pandemic poetry. The instructor was pedagogically stretched in translating what had always been an in-person research and poetry creation experience into a virtual setting. Still, the students responded with powerful poetic renditions. This article features the poetry of course members, along with a brief review of the course materials used and guidance provided to the class, and a methodological reflection is offered.

Cable Network Data Analytics System Cable Network Data Analytics System

Idilio Moncivais-Pinedo Patent Number: US-9497451-B2. Issued Nov 15, 2016
<https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/9497451>

A cable network data analytics system is configured to aggregate spectrum data sets of one or more readable devices connected to one or more cable networks. The spectrum data sets may include video spectrum data, which may indicate performance aspects of one or more standard channels of the cable network. The aggregated spectrum data may be analyzed against predetermined performance requirements. An alert may be generated if one or more performance aspects do not meet the predetermined performance requirements.

RECENT LECTURES AND PRESENTATIONS

SAP Sapphire 2024 Annual Conference
Topic: Predictive Analytics in Finance June 3-6, 2024
Audience: 15k attendees. Orange County Convention Center, Orlando, FL
<https://www.sap.com/events/sapphire.html?video=fec2b8b3-c17e-0010-bca6-c68f7e60039b>

Lumen Technologies 2024 Annual Leadership Conference
Topic: Improving Margin Management With Statistics Analysis March 13, 2024
Audience: 800 Attendees. Lumen Headquarters, Denver, CO
https://lumen.sharepoint.com/sites/Documents/Q124_Global_All_Hands/

Lumen Technologies 2024 Key Initiatives Conference
Topic: Advanced Analytics to Reduce Customer Churn January 24, 2024
Audience: 800 Attendees. Lumen Headquarters, Broomfield, CO
<https://lumen.sharepoint.com/sites/AdvancedServices/Documents/Jan24AllHands>

OTHER

Member of the University of Colorado at Colorado Springs' Artificial Intelligence Curriculum Advisory Board

Member of the Denver University's School of Engineering and Computer Science Industry Advisory Board

Associate Board Member of the Society of Cable Television Engineers

(SCTE) Member of CableLabs / OpenCable NDA and InGeNeOs®

standards group Member of the Colorado Association of Chiefs of Police

PERSONAL

Graduated from the “Comcast Emergent Leaders” 18-month long program on Sep/2008

Bilingual: English/Spanish.

LAW ENFORCEMENT

Lieutenant Commander with the Colorado Rangers Law Enforcement Shared Reserve (CLER), serving as the Director of Research since March/2023; Formerly a Captain in the same agency, appointed as the commanding officer of H Troop since Aug/2019. The rangers are a statewide government agency that supports 61 separate state, county, and local law enforcement agencies & governments across Colorado. It is the nation's first statewide shared Police Reserve.

Colorado Attorney General’s “Peace Officer Standards and Training” (POST) certified police officer with the Colorado Rangers Law Enforcement Shared Reserve (CLER). Currently serving as the Public Information Officer (PIO) and Director of Social Media for this government agency since Jun/2018.

Member of the CLER Firearms Instructor team. Holding current certifications as a POST Handgun Instructor, POST Patrol Rifle Instructor, Reflex Sight Instructor, and Axon Taser Instructor. In this capacity, I served as a “Full Skills” firearms instructor at the Flatrock Regional Training Center and Adams County Sheriff’s Office Police Academy, serving the northern region of the state.

Served with the CLER agency’s predecessor, the Colorado Mounted Rangers (CMR), since 2014 as a sworn ranger. Being the first ranger to serve as Director of state-wide operations with the rank of Major in charge of the newly created Operations Division from Jun/2016 to May/2017.