

JEREMY DOWDY

Data Scientist | Analytics Engineer

Prosper, TX | jeremydowdy.dev | linkedin.com/in/jeremy-dowdy | github.com/jeremydowdy | Jeremydowdy76@gmail.com | 903-474-4110

PROFESSIONAL SUMMARY

Highly technical Data Scientist and Analytics Engineer with a Master's in Advanced Data Analytics and 7+ years of progressive data platform experience at HOLT CAT. Specializes in enterprise lakehouse architecture, high-throughput ETL/ELT pipelines, predictive machine learning systems, automated pricing systems, and practical AI-assisted delivery. Known for translating industrial complexity into trusted data products that improve pricing, logistics, margin, and executive decision-making.

TECHNICAL SKILLS

Languages & Frameworks: Python (Pandas, NumPy, Scikit-learn, Seaborn, Matplotlib), SQL (window functions, CTEs), dbt, R, Excel
Data Platforms & Warehouses: Microsoft Fabric, Snowflake, Azure, Lakehouse Architecture, Medallion Layering (Bronze/Silver/Gold), PostgreSQL
Business Intelligence & Engineering: Power BI (DAX, semantic modeling), Tableau, Alteryx ETL automation, Streamlit
DevOps & AI Engineering: Azure DevOps, CI/CD pipelines, Git/GitHub, Cursor, LLM-assisted refactoring, agent workflows
Core Competencies: Dimensional modeling, predictive analytics, real-time pricing architecture, stakeholder discovery, API/webhook integration

PROFESSIONAL EXPERIENCE

HOLT CAT 7 yrs 2 mos

Data Scientist (Full-time) | Sep 2025 - Present | Dallas, TX (Remote)

- Architected and deployed a machine quoting and market-based pricing application, designing the data pipelines and algorithmic logic behind commercial machinery rate analysis and margin optimization.
- Spearheading "Project Galaxy," an end-to-end corporate data migration from legacy Snowflake, dbt, and Tableau footprints into a unified Microsoft Fabric Lakehouse environment.
- Productionized regression, classification, and K-Means clustering systems to predict industrial logistics bottlenecks and optimize fleet distribution networks.
- Pioneered AI-assisted development workflows using Cursor and LLM agents, accelerating refactoring and pipeline delivery speed by 35%+.
- Established reusable modeling and evaluation patterns for pricing and operations use cases, improving stakeholder trust in ML-assisted recommendations before production rollout.

Data Analyst 3 (Full-time) | Mar 2023 - Oct 2025 | Irving, TX

- Engineered scalable Azure and Snowflake ETL/ELT workflows and optimized warehouse indexing to reduce data pipeline ingestion latency by 25%.
- Conducted warehouse cost auditing, identifying redundant compute clusters and unoptimized queries to reduce cloud architecture overhead.
- Built production-grade Alteryx workflows to automate operational asset-health metrics, cutting manual reporting by 40%.
- Partnered directly with senior executive stakeholders to translate ambiguous commercial problems into programmatic technical requirements.
- Standardized KPI logic and reporting definitions across analytics outputs, reducing interpretation drift in executive and operational reviews.

Data Analyst 2 | May 2019 - Mar 2023 | Irving, TX

- Architected and managed enterprise Power BI and Tableau dashboard suites, providing visibility into billions of historical industrial transaction and telemetry records.
- Analyzed machinery transaction and telemetry records to isolate operational inefficiency patterns and support annual budget restructuring.
- Developed dimensional reporting logic and executive-facing views that connected raw operational activity to budget, utilization, and service planning decisions.

Fan Cloth 4 yrs 3 mos

Sales Executive | Nov 2014 - Jan 2019 | Texas

- Managed regional accounts and revenue targets through pipeline analysis, forecasting metrics, and logistics coordination.
- Built the business and operational foundation that now informs analytics discovery, product judgment, and stakeholder partnership.

KEY PROJECTS

Machine Quoting & Market-Based Pricing Engine

Architecture: Snowflake, Snowflake Notebooks, Streamlit, XGBoost, Python, SQL.

Architected a pricing intelligence system powered by classification and regression models to replace manual quoting spreadsheets with dynamic machine quoting recommendations for commercial teams, including feature engineering, scenario testing, and analyst-facing review workflows.

Enterprise Lakehouse Migration & Platform Modernization

Architecture: Microsoft Fabric, Power BI, dbt, Snowflake, SQL, Tableau.

Co-led migration patterns for moving legacy analytics assets into Fabric lakehouse and semantic model layers, standardizing metrics, reusable data contracts, and reporting expectations across business divisions.

Louie Boards Charcuterie CRM & Dashboard

Architecture: Next.js, Supabase, Vercel, Square API, TypeScript, PostgreSQL.

Built a custom CRM and order workflow for a boutique business, structuring customer history, order details, payment events, and fulfillment status into a durable data product with Square-driven workflow automation and a reporting foundation for repeat-order analysis.

Net Worth & Cashflow Management Engine

Architecture: Next.js, Supabase, PostgreSQL, Recharts, Vercel Crons.

Built a household finance platform with normalized transaction history, cashflow visibility, scheduled upkeep, caching, and sub-50ms query targets, creating a decision-ready layer for household planning rather than a simple ledger.

EDUCATION & CERTIFICATIONS

University of North Texas - Master of Science in Advanced Data Analytics, 2018 - 2019

East Texas A&M University - Bachelor of Science in Sports Management, 2010 - 2014

Microsoft Certified: Fabric Analytics Engineer Associate (DP-600)

Active participant in the dbt Global Community; attended Microsoft Fabric Community Conference (FabCon), 2026.