

Data and analytics engineering leader with 6+ years building data platforms, pipelines, embedded BI, and self-service analytics at high-growth startups. Currently scaling a data team and using dbt, Snowflake, Dagster, and Looker to solve painful problems for internal teams and customers.

PROFESSIONAL EXPERIENCE

Regal.ai | New York, NY

February 2023–Present

Manager, Data & Analytics Engineering (Aug 2024 – Present) | Senior Data Analyst (Feb 2023 – Aug 2024)

- **Team Leadership and Data Platform Ownership:** Built and scaled data function. Hired and developed direct report through promotion to Data Scientist, and managed workstreams for two data engineers. Own end-to-end pipeline (S3 → dbt → Looker) serving near real-time dashboards to 200+ customers. Refactored dbt and LookML repositories with industry-standard patterns and enabled cross-functional teams to self-serve, removing the data team as a bottleneck.
- **LookML & Dashboard Automation Platform:** Designed and built a fully automated, metadata-driven system that eliminates manual LookML and dashboard work for customer-configured AI Agent metrics across multiple brands (~5,000 lines of production code, orchestrated by Dagster, running hourly). The LookML layer syncs agent metadata from Snowflake, auto-generates dimension files, commits to GitHub via PR with Slack alerts, and delivers self-service analytics within hours of customer configuration. The dashboard layer programmatically creates and maintains 1,500+ unique Looker dashboards (10,000+ tiles)—provisioning brand-scoped folders, generating dashboards with deterministic SHA-256-based slugs (matching frontend embedding logic), and building channel-aware tile layouts (voice vs. SMS) with appropriate chart types (single-value KPIs, time-series columns, stacked bar, scatter, pie, and grid tables). Implements an incremental processing strategy that pre-loads tracking data into in-memory DataFrames (4 bulk Snowflake queries instead of ~6,000 per-agent queries), performs anti-join logic to process only agents with new or missing dashboards/tiles, and skips 80–90% of agents on typical runs—reducing Looker API load and runtime significantly. Handles scale and resilience with brand-level resource reuse (single SDK instance, cached folder lookups, shared explore metadata), bulk tile metadata flush with per-row fallback, slug conflict resolution for soft-deleted dashboards, dry-run mode for safe testing, and structured per-agent error isolation so one failure doesn't block the remaining 1,499 agents. Reads from dbt-modeled dimensional tables, writes tracking state back to Snowflake, and supports both single-agent and multi-agent execution modes via Dagster config.
Tech: Python, Dagster, Looker SDK/API, Snowflake, Pandas, LookML, SHA-256 slug generation
- **Regal IQ–AI Agent Knowledge Base Analysis:** Launched the company's first data product, surfacing gaps in AI voice agent knowledge bases and prompt design across thousands of conversations. Built dbt transcript models and LLM/ML pipeline (topic extraction → embeddings → HDBSCAN clustering → LLM labeling) orchestrated via Dagster, delivering insights to 200+ customers through embedded BI.
- **Data Warehouse Refactor & Architecture:** Reduced production runtime by 26% (55 to 40 minutes) and cut Snowflake spend by \$60K (50%) by refactoring models to incrementally process heavy event data, clustering tables, and implementing Kimball methodology. Improved developer experience through dbt best practices and slim CI/CD testing in GitHub Actions. Provided compute cost visibility in Looker by modeling the select.dev package.
- **Latency Reduction Analysis:** Improved progressive dialer call routing by 60%, reducing contact abandon rate, by building a data model tracking event timestamps through the call journey with millisecond precision to identify latency bottlenecks.
- **Profile Key Migration:** Transitioned primary user identifier in data warehouse from phone number to profile ID across 300+ dbt models, achieving parity between application and data products. Cloned dbt repo to enable gradual migration, wrote validation scripts comparing production data to new tables for core KPIs, and used LookML liquid variables for blue-green deployment.

Navan (TripActions) | Palo Alto, CA & New York, NY

May 2019–February 2023

Data Analyst (Oct 2021 – Feb 2023) | Manager, Sales Operations (May 2021 – Sep 2021) | Senior Sales Operations Analyst (Apr 2020 – Apr 2021) | Sales Operations Analyst (May 2019 – Mar 2020)

- **Marketing Attribution:** Doubled personal travel bookings by building the first email-to-booking attribution model, tying campaigns directly to conversions and providing visibility into marketing ROI.
- **A/B Testing:** Enabled engineering to reliably roll out new features by standardizing experimentation. Built a dynamic dbt model that auto-ingests experiments and calculates Z-scores, confidence intervals, and effect sizes in Tableau.
- **Territory Re-Segmentation & Account Score Model:** Fully automated territory assignment for 75+ AEs across 580K accounts and 4 segments. Built predictive scoring models in dbt for TAM calculations and auto-assignment, reverse-ETL'd to Salesforce via Census, giving sales visibility into scoring inputs.

SKILLS & CERTIFICATIONS

SQL, dbt, Dagster, JinjaSQL, Git, Python, Snowflake, LookML, Tableau, Product Analytics, Data Modeling (Kimball), Revenue Analytics, Salesforce, Data Visualization, Salesforce Certified Administrator

EDUCATION

The University of Massachusetts Amherst | *BA Economics & Minor Political Science* |
Amherst, MA

December 2016