

ZACHARY SIMS

zsims413@gmail.com

(678) 637-9036

<https://zsims413.com>

<https://github.com/zack-sims413>

SKILLS

Languages & Cloud Architecture: Python, SQL, Snowflake, AWS (Lambda, S3, Redshift), Google Cloud Platform (BigQuery, Cloud Functions), GitHub CI/CD.

Machine Learning & Applied AI: Scikit-learn, XGBoost, Hyperparameter Tuning (Grid/Randomized Search), Text Embeddings & Semantic Search, Snowflake Cortex, RAG Architectures, LLM Inference (Groq, OpenAI gpt-4o-mini).

Data Engineering & Pipelines: Custom API Integrations, ETL Frameworks, dlt, dbt, Supabase (PostgreSQL, pgvector, PostGIS / Spatial Data), DuckDB (MotherDuck).

Analytics & Visualization: Tableau, Power BI, Streamlit, Evidence.dev.

PROFESSIONAL EXPERIENCE

Atlanta Journal-Constitution – Atlanta, GA

Sr. Manager – Advanced Analytics & AI, September 2025 – Present

- Oversaw the end-to-end development and deployment of an in-house content recommendation model for personalized newsletters, driving a 50% increase in click-through rates compared to the legacy third-party system.
- Engineered an automated ingestion pipeline to generate text embeddings for new articles, enabling semantic mapping and clustering of similar content.
- Built and deployed lifecycle-specific predictive models to assess churn propensity among subscribers, empowering the retention team with proactive, data-driven intervention strategies.
- Co-architected an end-to-end dynamic paywall system, deploying an AWS Lambda-based inference engine that processes real-time events alongside hourly pre-aggregated user & content features retrieved from a Valkey cache.
- Designed automated feature pipelines moving data from Snowflake to S3 and Valkey to support model serving. Executed a targeted data collection experiment on 10% of traffic to generate training data, integrating real-time meter assignments with Amplitude for closed-loop product analytics.
- Conducted 30-day subscriber onboarding analyses to identify key relationships between early platform engagement and long-term retention likelihood.

Manager – Analytics & Data, November 2023 – September 2025

- Led the migration of legacy data & processes from Amazon Redshift to Snowflake, optimizing DDL, stored procedures, UDFs, and Lambda functions.
- Developed and deployed unsupervised clustering models to segment users based on content consumption and engagement patterns.
- Designed a chatbot prototype using Streamlit & Snowflake Cortex, allowing non-technical stakeholders to interact with data via natural language.
- Established a GitHub-based development framework to improve code quality and deployment efficiency for the Analytics & BI teams.
- Rebuilt a centralized customer profile, providing a 360-degree view of the user, which was exported to external systems for targeting and segmentation.

Advanced Analytics Developer, February 2022 – November 2023

- Developed and deployed an audience engagement classification model using AWS SageMaker to segment subscribers based on behavioral patterns.
- Automated sentiment & emotion analysis of articles using pre-trained Hugging Face models on AWS SageMaker, with inference pipelines updating predictions in Redshift.
- Designed and built ETL pipelines using AWS Lambda & APIs to ingest third-party data (Facebook, Twitter, App Figures, Megaphone, etc.).

- Created automated web scrapers to track competitor headlines and social media trends.

SpendHQ – Peachtree Corners, GA

Associate, Business Intelligence & Data Science, June 2021 – February 2022

- Developed Alteryx workflows to ingest and clean multi-source spend data for client analytics.
- Integrated Python-based automation with Alteryx's API to optimize ETL processes.

Cox Media Group – Atlanta, GA

Senior Data Analyst, November 2017 – June 2021

- Built a forecasting model using Google BigQuery to set annual traffic goals for CMG's digital platforms.
- Developed classification models on Google Cloud to categorize push notifications and improve messaging strategies.

Turner Broadcasting System, Inc. – Atlanta, GA

Analyst TBS/TNT Strategic Research, October 2016 – November 2017

Competitive Sports Analysis – Atlanta, GA

Analytics Intern, April 2016 – October 2016

PERSONAL PROJECTS

WiseBites – Gluten-Free Restaurant Discovery Platform

- Architected an end-to-end, AI-powered restaurant discovery platform tailored for users with gluten restrictions and celiac disease, enabling rapid discovery across standard search and map-based interfaces.
- Integrated Groq (LLaMA 3.3) and Python (FastAPI) to process and deliver low-latency AI summarizations of user reviews, highlighting gluten-free safety and dining experiences.
- Developed a prototype conversational AI agent ("Wisey") utilizing a hybrid vector search (RAG) architecture via Supabase pgvector and OpenAI (gpt-4o-mini) to extract user intent and enable semantic matching over exact keyword searches.
- Designed a modern ELT pipeline utilizing dlt to extract operational data into a serverless OLAP warehouse (MotherDuck/DuckDB), transforming raw data via dbt to optimize analytical querying.
- Managed robust API integrations with Google Places and SerpApi for the continuous ingestion of dynamic spatial data and user reviews to feed the AI models.
- Implemented a "BI as code" reporting layer using Evidence.dev to generate version-controlled, static dashboards for tracking user engagement and system performance.
- Engineered an event-driven user lifecycle pipeline, capturing transactional data via Stripe webhooks to synchronize user states within Supabase, automatically triggering targeted onboarding and engagement workflows via Brevo.

EDUCATION

Georgia Institute of Technology - Atlanta, GA

Master of Science in Analytics, Jan 2019-Dec 2022

The University of Georgia - Athens, GA

Bachelor of Arts in Journalism, May 2016

Bachelor of Science in Statistics, May 2016

Certificate Program - Sports Media

NCAA Division I Student-Athlete (Track & Cross Country)