

Chris Samanant

Twin Cities, MN • (612) 720-8483 • chris.samanant.jobs@gmail.com • christopher.samanant@optum.com

EDUCATION

UNIVERSITY OF MINNESOTA, Minneapolis, MN

Carlson School of Management

Master of Science in Business Analytics

May 2020

GPA 3.3 | MSBA Fellowship Scholarship | MSBA Ambassador

UNIVERSITY OF MINNESOTA, Minneapolis, MN

Bachelor of Science in Economics

May 2019

GPA 3.29

EXPERIENCE

INDEPENDENT, HeyBye.ai (Website: Search Engineering Platform)

Full Stack Data Scientist

Mar 2025 – Present

- Built a distributed framework to iterate daily over a public search API and process all existing and new documents in parallelized GKE nodes. Each record is automatically enriched with OpenAI-powered generative pipelines and embeddings, then upserted into a self-managed Vespa vector database on GCP for faster and more directed information retrieval.
- Engineered a consensus-driven evaluation stack with AI relevance metrics and a UI/API voting system that aggregates human ratings to gain statistical confidence in relevance labels. These confidence-weighted labels enable automated, comparable evaluations of different search algorithms using precision@k, recall@k, MAP, and NDCG.
- Developed three additional production-ready applications showcasing vector search application management:
 - Intro Search & Collective History – word-cloud visualization of collective search history with freeform user search with results linking directly to online publication.
 - Search & Relevance Demo – interactive freeform search and labeling interface displaying real-time precision/recall/MAP/NDCG metrics at k = 3, 5, 10.
 - Fine-Tuning Sandbox – before/after freeform re-ranking Demonstration using localized keyphrase embeddings and polynomial & coefficient weight calculations.
- Built and deployed HeyBye.ai, a fully interactive Next.js / React platform unifying all project components into a cohesive live experience. The site embeds every backend system — from Vespa vector search to real-time evaluation and fine-tuning — within a seamless front-end flow. Features include guided voice-over walkthroughs, integrated API visualizations, real-time model comparisons, and a contact portal with résumé download and Calendly scheduling. Designed as both a user-facing demonstration and a full production interface for vector-based search R&D.
- Deployed components are deployed within a unified GCP VPC. The Vespa vector cluster and integrated applications—voter, search/relevance, fine-tuning, and the showcase site HeyBye.ai run on autoscaling GKE workloads. Daily ingestion and enrichment pipelines scale across dynamic node pools based on input volume, with operational visibility through Grafana and Splunk.
- Stack: GCP · OpenAI · React/Next.js · Vespa Vector DB · PostgreSQL · FastAPI · Python · Docker · Github Actions · Grafana · Statistics

OPTUM, Eden Prairie, MN

Senior Data Scientist

Jan 2022 – Mar 2025

- Part of proof-of-concept team seeking to improve Benefits Search using ML research.
- Wrote Python scripts to clean data for exploration using Github Actions. The resultant data and analysis were initially transferred to our S3 bucket automatically, with later versions instituting temporary parallelized servers to convert and transfer augmented data directly to our vector databases for search applications.
- Designed a Vespa Application for this data, complete with multi-vector NLP search. Iterated from local, single container development to cloud deployments in Azure and GCP with test, stage, and prod environments across several availability zones.
- Wrote in Python to comprehensively test a variety of search algorithms and NLP models to optimize search accuracy, saving results to a database and using APIs to perform side by side comparisons on a Streamlit UI and select the most appropriate model.
- Worked in conjunction with AI/ML Engineers to combine an OpenAI generative model to act as a question answering assistant, which would parse the most relevant documents to efficiently provide direct answers.

Data Engineer – Technology Development Associate

June 2021 – Jan 2022

- Validated big data migrations (50+ million rows and thousands of variables) between MS SQL Server and Snowflake using SSIS and Python. Coordinated findings to improve data quality and analysis.
- Programmed in SQL to fix and improve efficiency on multiple stored procedures that run regularly.
- Created multiple monitors that would alert the business when data meeting certain conditions is added to the dataset.

CARLSON ANALYTICS LAB, Minneapolis, MN

Analytical Student Consultant

June 2019 – May 2021

Client: Hospitality and Entertainment Organization (Experiential Learning Project)

- Analyzed client's new revenue stream across different segments and time periods and delivered descriptive analytics presentation to client with clear marketing response strategy suggestions.
- Constructed a causal analytics framework for the new data source to clarify the effect of the new data on existing revenue streams.

Client: Mall of America (Exploratory Analytics)

- Applied analytical techniques using R on over three years of unexplored call log data to make predictions to serve to optimize operations regarding staffing scheduling and general procedure.
- Volunteered to lead a data analytics team's final client presentation to communicate findings and suggestions.

Data Science Projects

Image Classification Across Platforms: Compared performance and efficiency of image classification across three platforms – Jupyter Notebook, PySpark, and Amazon Sagemaker. Our team of four found Amazon Sagemaker to perform best across all categories for datasets up to 16K GB, including categories such as ease of use, accuracy, and speed. We then presented these findings to industry professionals at a technology presentation convention.

Uplift Target Marketing: Assumed a leading team role to showcase, promote, and demonstrate practicality of a Python package's ability to leverage A/B testing datasets to automatically design a model and identify customers to target with marketing and advertising campaigns based on identifiable demographics. Highlighted the packages applicability to market to any industry's customer base. Presented to industry professionals at a technology presentation convention.

SKILLS

Programming Experience: Anaconda, RStudio, MySQL, Tableau, MATLAB, Mathematica, AutoCAD, IntelliJ

Computer Language Experience: SQL, Python, R, C++, C, Java, JavaScript

Updated Proficiencies 2025: Vespa, Github, Postman, MongoDB, Postgres, Flask, Streamlit, HuggingFace, OpenAI, Docker, Kubernetes, Helm, Azure, GCP

REFERENCES

Optum Key references:

Current Supervisor Name: Ross Rosemark | Email: ross.rosemark@optum.com | 612-715-6397

Former Optum Supervisor: Andrei Filimonov | Email: afilemonov@gmail.com | (612) 269-1343