

An AI (Artificial Intelligence)-Powered Approach to Accessing Catchment Area Data

D. Stone, T. Onega

Huntsman Cancer Institute at the University of Utah

1. Background

Cancer centers are expected to identify disparities, guide outreach, and demonstrate community impact. Data required for these analyses—drawn from registries, electronic health records, surveys, and other datasets—are stored in complex relational databases requiring technical expertise to access. As a result, program staff rely on analysts for queries, which limits timeliness, flexibility, and integration of data into day-to-day decision making. Existing dashboards address predefined sets of questions and require lots of time to evolve to accommodate growing needs. This gap between data availability and accessibility constrains how effectively cancer centers can use catchment analytics to inform equity-driven interventions and meet cancer center reporting needs.

2. Goals

This project aims to improve accessibility and efficiency of catchment area data usage while maintaining accuracy. Objectives include enabling users to obtain catchment area data through natural language queries and reducing analyst time spent on ad hoc requests. Evaluation metrics include routing of queries to appropriate logic branch of the agent, Structured Query Language (SQL) generation accuracy, and response satisfaction. Success is defined as taking a natural language query and accurately retrieving relevant data.

3. Solution and Methods

We developed a pilot AI-powered interface for our catchment area database. The system allows users to submit plain language questions that are automatically translated into SQL, executed on our Microsoft SQL Server database, and returned as structured tables, or if not answerable with a query, an appropriate response is returned. The application uses a modular LangGraph workflow integrating a locally hosted large language model with a retrieval layer that indexes database schemas. These resources ground the model in accurate table and field definitions during routing, query generation, and response generation. Model performance was tested using two query sets: historical requests previously submitted to our team and prompts created by analysts and staff to probe accuracy, routing, and robustness.

4. Outcomes

Preliminary testing demonstrates moderate performance. The system routed 96.7 percent of queries to the appropriate logic branch. Analyst scoring of generated SQL averaged 2.7/5, reflecting a mixture of valid and invalid queries. User-facing responses scored 4.2/5, indicating strong clarity and interpretability for non-technical users. Although a time savings analysis has not yet been completed, the system produces results within seconds, representing a substantial reduction relative to analysts' mediated turnaround times.

5. Lessons Learned and Future Directions

Early testing shows that a natural language interface has potential to expand access to catchment area data for both technical and non-technical users. The agent still requires improvement of accuracy to reach the standards of performance for a reliable AI agent. Future work will focus on expanding documentation for richer context, enabling visualizations, developing security guardrails, and conducting broader user testing with program staff. Continued evaluation of accuracy and time savings will guide development toward a scalable, secure, and equity-promoting tool for catchment area analytics.