Eurographics Workshop on Visual Computing for Biology and Medicine (2025)
T. Gerrits and C. Gillmann (Poster Chairs)

Poster

AI-Powered Dashboards for Cohort Study Monitoring and Analysis

H. Maharlou[1] and S. Oeltze-Jafra[1]

[1]Peter L. Reichertz Institute, Hannover Medical School, Germany

Abstract

We present an AI-powered dashboard system designed to improve the lifecycle of cohort studies by offering intelligent support during both data collection and post-study analysis. The dashboard combines natural language interfaces with generative agents to assist in data monitoring, chart creation, anomaly detection, and result interpretation. This paper outlines a modular architecture for flexible and scalable deployment and demonstrates its potential through early-stage prototypes developed for medical research.

CCS Concepts

• ***Applied computing*** → ***Life and medical sciences;*** • ***Computing methodologies*** → ***Artificial intelligence;***

1. Introduction

Researchers managing large-scale cohort studies often face operational bottlenecks, such as delays in identifying missing assessments or challenges in comparing recruitment across multiple sites. Addressing these issues typically requires support from technical staff, given the complexity of structured database queries [JNQ11].

To reduce these barriers, we are developing *CohortCanvas*—an AI-augmented dashboard that supports both real-time monitoring during data acquisition and flexible post-study analysis. It enables users to pose high-level questions (e.g., "Show recruitment trends over the last 3 months by site") and receive visual results without writing any code.

Key features:

- **Monitoring phase:** Real-time participant tracking, protocol adherence, and detection of missing or anomalous data [HWS*24].
- **Analysis phase:** Natural-language-driven filtering, subgroup comparison, and automated chart generation.

CohortCanvas leverages large language models (LLMs), retrieval-augmented generation (RAG), and semantic embeddings to lower the technical threshold for working with cohort data. These techniques translate user queries into structured operations over cohort datasets [XZZZ25], improving accessibility across disciplines and supporting more inclusive, real-time data exploration.

2. Methods

CohortCanvas is designed as a modular system composed of independent agents, each responsible for a specific aspect of user interaction, data access, or visualization. The system aims to reduce the technical complexity of interacting with cohort data while preserving flexibility and security.

Orchestration and Agent Architecture: At the core of the platform is an orchestration layer that receives natural language requests from the user and determines the appropriate agent to handle the task. For instance, a request to visualize a metric over time is routed to the chart generation agent, while a request to find variable names is handled by the variable finder agent.

Semantic Variable Finder: The variable finder agent maps natural language phrases to dataset variables using a semantic embedding pipeline. User queries and variable metadata (from the data dictionary) are encoded using the `sentence-transformers` library and stored in a PostgreSQL database with PGVector support. During a query, the system retrieves the top matches based on cosine similarity. This allows flexible queries such as "anxiety score" to identify variables like `GAD7_total` — the Generalized Anxiety Disorder (GAD-7) scale, even if the wording differs.

SQL and Chart Generator: Once the relevant variables are identified, a chart generation agent uses a locally hosted LLM (LLaMA3 via Ollama) to generate SQL queries and Vega-Lite specifications. Prompts are tightly structured to reduce hallucination and ensure syntactic validity. This enables automatic chart generation for user questions like "Compare recruitment trends by site over the last three months."

Interactive Dashboard Interface: The user interface is developed in Next.js and uses the React Grid Layout library to support a modular, section-based design. Users can create, resize, and arrange individual chart blocks on the canvas. Each block can independently connect to the AI assistant, enabling localized requests such as filtering by age group or comparing outcomes across subgroups. Linked blocks share filters or selections, allowing synchronized updates across charts for comparative analysis.

Backend and Deployment: Data is stored in PostgreSQL and managed through Prisma ORM, which simplifies schema handling and enforces type safety. The system is containerized with Docker, supporting local and cloud deployment. Authentication and role-based access control are provided via NextAuth to ensure secure multi-user access.

Example Workflow: During the monitoring phase of a study, a site coordinator might ask, “How is participant recruitment progressing by month at each hospital?” The system interprets the query, identifies the necessary variables, constructs the SQL, and displays a chart—all without requiring manual query writing. In the analysis phase, a researcher might request a filtered dataset of participants over 60 years old with moderate depression to visualize treatment adherence over time. The AI agents facilitate this interaction end-to-end, from variable selection to visualization.

2.1 System Architecture

Figure 1 gives a general view of how CohortCanvas is set up. The system runs inside Docker containers and consists of several services that communicate through REST APIs.

Users access the dashboard through a web browser over HTTPS. The frontend is built with Next.js and connects to a local PostgreSQL database, which takes care of layout saving and user session tracking. We use NextAuth for login and role management. Prisma handles the database schema, streamlining query construction and reducing the need for manual SQL scripting.

In the backend, a lightweight orchestration service processes natural language input from the user and determines the appropriate action. Depending on the identified intent—such as variable retrieval, SQL generation, or chart construction—the request is routed to the corresponding AI agent. Each agent operates under constrained prompt templates to ensure the generation of syntactically valid SQL and renderable Vega-Lite chart specifications.

For integration with external data sources, the system employs a schema crawler and pre-computed vector embeddings of variable metadata to enable semantic matching between user queries and dataset variables (e.g., GAD-7 for anxiety score).

As a data security measure, all AI-generated SQL statements are passed through a guard node prior to execution. The guard enforces a read-only policy, permitting only commands such as SELECT, thereby preventing accidental or unauthorized data modification.

For system monitoring, we use Loki and Grafana to collect logs and metrics. Each container includes a Promtail sidecar that forwards logs to the stack for real-time tracking.

Overall, the system is modular, secure, and designed for explainability and live interaction — especially useful in multi-site studies with distributed data sources.

3. Results and Use Case

CohortCanvas is currently in early-stage development but has shown promising functionality across the following areas:

- **Monitoring phase**: Enables real-time visualization of recruitment trends, detection of missing data, and tracking of protocol adherence across study sites.

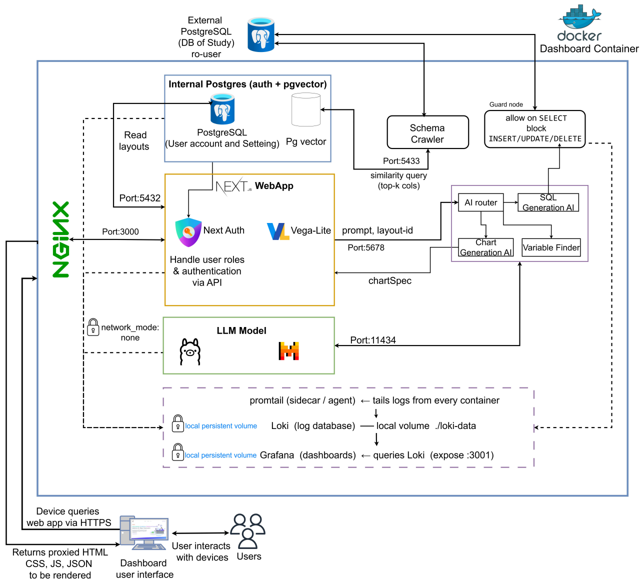

Figure 1: *System architecture of the CohortCanvas dashboard, showing interactions between user interface, authentication, AI agents, databases, and monitoring tools.*

- **Analysis phase**: Researchers can create custom charts, filter subgroups, and export datasets using natural language.
- **Usability**: Non-technical users (e.g., clinicians) can interact with study data independently, reducing reliance on database experts.

In a preliminary evaluation involving 10 representative queries, the variable finder agent achieved **82% top-1 match accuracy**—i.e., the correct variable was returned as the single best match. For successful cases, the system generated valid Vega-Lite chart specifications in under 5 seconds.

4. Discussion

CohortCanvas shows how AI agents can enhance interactivity and accessibility in clinical dashboards. By combining natural language processing, variable search, and chart generation, it lowers technical barriers for researchers. Its modular design supports flexible deployment across diverse cohort infrastructures. While early results are promising, further evaluation with larger datasets and real users is required. Building on our earlier P4D Dashboard [MKM*25], which supports both reactive and proactive detection of suspicious study data and biases, we will extend CohortCanvas accordingly as part of future work, with emphasis on stronger grounding, multimodal data integration, and reproducibility in real-world studies.

Acknowledgments

This work was carried out in the context of the P4D study (German Clinical Trial Register DRKS00032215), whose support is gratefully acknowledged.

References

[HWS*24] Honushefsky A., Wagner E. S., Sheridan K., Spickard K. M., LeMasters W. R., Walter C. N., Beaver T., Lennon A. M., Papadopoulos N., Rahm A. K., Buchanan A. H.: Real-time evaluation and adaptation to facilitate rapid recruitment in a large, prospective cohort study. *BMC Health Services Research 24* (12 2024). doi:10.1186/s12913-024-10750-5. 1

[JNQ11] Jagadish H., Nandi A., Qian L.: Organic databases. In *International Workshop on Databases in Networked Information Systems* (2011), Springer, pp. 49–63. 1

[MKM*25] Maharlou H., Krause E., Müntefering F., Voges J., Weihs A., Lucht M., Grabe H. J., Pollmann I., Mueller F.-J., Weber H., Oeltze-Jafra S.: The p4d dashboard: A platform for monitoring clinical studies. *Stud Health Technol Inform 316* (2025), 495–499. doi:10.3233/SHTI250889. 2

[XZZZ25] Xiang L., Zhao Y., Zhang Y., Zong C.: A survey of large language models in discipline-specific research: Challenges, methods and opportunities. *Studies in Informatics and Control* (7 2025). URL: http://arxiv.org/abs/2507.08425http://dx.doi.org/10.24846/v34i1y202501, doi:10.24846/v34i1y202501. 1