

IntegratedViewer (IVe) — Investigator's Explorer: In-context, SQL-Free, Natural-Language Exploration of Multimodal Critical-Care Data in the OMOP CDM

Delgersuren Bold¹, Eric Rosenthal², Xiao Hu¹, on behalf of the CHoRUS Consortium

¹ Center for Data Science, Emory University, Atlanta, GA; ² Harvard University, Massachusetts General Hospital, Boston, MA

Background

Our goal is multimodal clinical event annotation and adjudication: having domain experts mark and agree on clinical events using the full patient picture, including structured records, waveforms, and images. Expert-adjudicated labels are the ground truth needed to train and test models, and they are especially valuable for critical-care datasets such as CHoRUS Bridge2AI, which adds waveforms and medical images to the usual OMOP tables.¹ But this work is hard today. The OMOP Common Data Model (CDM) gives the OHDSI community a shared way to store data, yet working with a live CDM still requires SQL and deep knowledge of its many tables. A clinician who wants to review patients with sepsis, or see who received norepinephrine, must first turn that idea into queries. Tools like ATLAS and ATHENA help with cohorts and vocabularies, but none lets an expert search concepts, browse cohorts, read a record, view waveforms, and annotate events in one place. To make annotation possible, we first needed two things: a collection of endpoints that can be annotated, and easy, organized layouts and widgets to show the right data for each task. We built IntegratedViewer (IVe), the Investigator's Explorer, to provide them. IVe is an open-source web app that runs inside the CHoRUS Bridge2AI secure enclave and lets users go from a clinical idea straight to the data, without writing SQL.

Methods

IVe is a web application (React, TypeScript, Vite, Redux, MUI) with a Node.js/Express backend. Every OMOP table is served by one simple, well-documented REST API (Figure 1). Each table has a GET route to read rows and a POST /search route to filter them. This collection of endpoints is the unit we annotate, and it powers two key features. First, the Workspace lets users build dashboards by dragging in widgets; each widget is backed by a single endpoint, so an expert can organize the exact views needed for an annotation task and save them as a layout, without coding. Second, the Clinical Agent uses the same endpoints as its context: we collect the /search endpoints (for example, /omop/person/search) and give them to the agent, so it can turn a plain-English question into the right queries. The model never sees raw patient data, only the API. IVe has five main areas: (1) Concept Search across SNOMED, RxNorm, and ICD-10 (about 9.5M concepts); (2) Cohort Explorer, to browse cohort definitions and open one subject's full record; (3) Workspace, for the annotation layouts; (4) API Endpoints, to browse and run every call; and (5) the Clinical Agent. A waveform viewer shows physiologic signals (ECG, pulse, and breathing) so events can be reviewed against the bedside trace. New tables and data sources load automatically, so the endpoint collection is easy to extend. Logins use secure tokens with role-based access, and the app works with different databases (SQLite in development, PostgreSQL in production). IVe runs on the CHoRUS Bridge2AI enclave over the first 50,000 ICU admissions (OMOP CDM v5.4). Source code: <https://github.com/chorus-ai/IVe>

Figure 1. The OMOP REST API (CHoRUS production server). Each table has a GET route to read rows and a POST /search route to filter them. This annotatable collection of endpoints powers the Workspace widgets (one endpoint per widget) and gives the Clinical Agent its search context.


Servers		https://mgh-chorus.eastus2.cloudapp.azure.com/api - production server	Authorize
OMOP OMOP CDM clinical data			
GET	/omop/person	Get all persons	
POST	/omop/person/search	Search persons	
GET	/omop/person/{pid}	Get person by ID	
GET	/omop/cohort	Get all cohorts	
GET	/omop/cohort/definition/{cohort_definition_id}	Get cohort by definition ID	
POST	/omop/cohort/search	Search cohorts	
GET	/omop/cohort_definition	Get all cohort definitions	
GET	/omop/cohort_definition/{cohort_definition_id}	Get cohort definition by ID	
POST	/omop/cohort_definition/search	Search cohort definitions	
GET	/omop/condition_occurrence	Get all condition occurrences	
GET	/omop/condition_occurrence/person/{person_id}	Get condition occurrences by person	
GET	/omop/condition_occurrence/visit/{visit_occurrence_id}	Get condition occurrences by visit	
POST	/omop/condition_occurrence/search	Search condition occurrences	
GET	/omop/death	Get all deaths	
GET	/omop/death/person/{person_id}	Get deaths by person	

Results

I've runs on the live CHoRUS Bridge2AI data and shows every OMOP table with its current row count, plus image and waveform media. A user types a clinical term in Search and lands on the matching table with that concept already applied as a filter. Opening a cohort subject shows their full record: demographics, counts for each domain, and an interactive timeline (Figure 2a). For example, Subject #4009 has 4,927 events, including 21,278 measurements, over a 744-hour ICU stay, with the top conditions, drugs, and procedures listed by standard term — the events experts will annotate. Any filter can be saved as a reusable, taggable endpoint, which becomes a target for annotation. In the Workspace, experts build annotation layouts from these endpoints and save them for reuse. One ICU example (Figure 3a) shows predicted organ-failure (SOFA) trends, a risk chart comparing now versus six hours ahead, key labs, and suggested next steps. The Clinical Agent answers a request like “Summarize Person #18462” by running four OMOP queries and returning a clear result (Figure 2b). Because CHoRUS is multimodal, the waveform viewer streams a subject’s ECG, pulse, and breathing signals and lets the user scrub through them (Figure 3b), so an event can be confirmed against the trace. All of this stays in one app, linked back to the OMOP record. With the annotatable endpoints and layouts now in place, the next step is to have domain experts generate event annotations by domain, and then to test model-generated annotations against this expert ground truth.

Figure 2. Single-subject review and annotation. (a) Cohort subject view: demographics, per-domain counts, and an interactive events timeline (Subject #4009: 4,927 events over a 744-hour ICU stay) with ranked top conditions, drugs, and procedures — the events experts will annotate. (b) Clinical Agent: a plain-English request (“Summarize Person #18462”) answered by running four OMOP queries.

(a) Cohort subject view



(b) Clinical Agent






Figure 3. Multimodal dashboards and signals. (a) A saved Workspace annotation layout (“ICU Deterioration Forecast, 2–6 h”): SOFA trends, a risk chart (now vs. +6 h), key labs, and suggested next steps. (b) The waveform viewer: ECG (leads II/V/aVR), pulse, and breathing, opened from a subject and scrolled to confirm events against the trace.

(a) Workspace annotation layout



(b) Waveform viewer



Conclusion

Our aim is multimodal clinical event annotation and adjudication, and IVE is the foundation for it. In one app, experts can search concepts, study cohorts, open patient records, view waveforms, and ask questions in plain English, without SQL. A single, annotatable collection of endpoints powers both the dashboards and the agent, which keeps the app simple and easy to extend. Running on 50,000 ICU admissions in the CHoRUS Bridge2AI enclave, IVE brings structured records and waveforms into one view so events can be marked in context. Next, domain experts will create annotations by domain, and we will compare model-generated annotations against this ground truth. IVE is open source, and we welcome the community to try it and to add new widgets and data sources.

References

1. Rosenthal E, Clermont G, Williams A, et al. 1713: The initial 50,000 ICU admissions within the CHoRUS Bridge2AI multimodal dataset. Crit Care Med. 2026;54(3S). doi:10.1097/01.ccm.0001188848.23432.3c.