AMIA

INFORMATICS PROFESSIONALS. LEADING THE WAY.

2026 Annual Symposium

#AMIA2026



A SMART-on-FHIR Knowledge-Graph App for Point-of-Care Adverse Drug Reaction Exploration and Prediction

Session Code: S73

David Guo, PhD

Department of Information Systems and Cyber Security
College of AI, Cyber, and Computing
The University of Texas at San Antonio

Connect with me @ www.nexusmed.ai 👉

Disclosure of Relevant Financial Relationships

The following presenter(s) have relevant financial relationship(s) with ineligible companies to disclose.

David Guo, PhD – the founder of Nexusmed.ai

* All of the relevant financial relationships listed for these individual(s) have been mitigated.

THE PROBLEM

The evidence exists. It just cannot be reached.

1.25M
serious ADR cases reported to FAERS in 2022

~175K
fatalities among those reports

3–7%
of hospital admissions are ADR-related

Much of this harm is preventable given the post-marketing evidence sitting in FDA's FAERS DB

Disclaimer: FAERS should not be used in isolation by prescribers. It is designed for regulatory analysts, not for the 15-minute clinical encounter in which treatment decisions must actually be made.

WHY IT STAYS UNREACHABLE

Three Barriers between FAERS and the Bedside

STRUCTURE

Drug names are uncontrolled free text

“LASIX” · “furosemide” · “FUROSEMIDE 40MG TAB”

A naive query for one drug misses most of its own cases.

CONTEXT

FAERS answers the wrong question

It answers “has this event been reported for this drug?”

A prescriber needs: “…for a 78-year-old woman, at this weight, for this indication?” Public dashboards cannot express that.

WORKFLOW

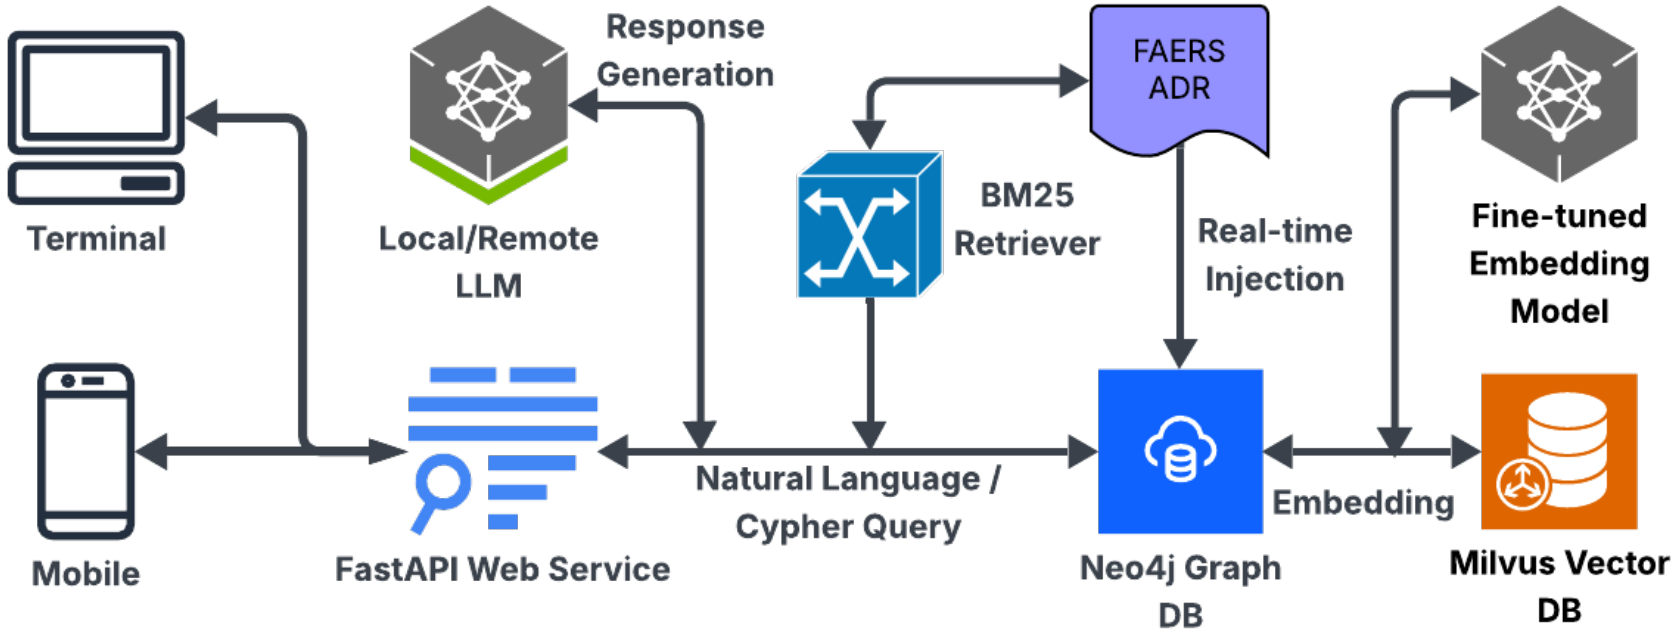
The tool sits outside the EHR

Even an expert must leave the chart and re-key age, sex, weight, drug list and indication into a separate system.

No time for the detour due to the typical 15-minute encounter limit

APPROACH

Neither Knowledge Graph Nor Vector Index Suffices Alone



A vector store alone cannot audit. A graph alone cannot handle messy drug nomenclature at scale.

Neo4j holds normalized relational structure · Milvus holds the semantic index · retrieval crosses between them. Every component is open-source and runs on-premises, including LLM inference.

KNOWLEDGE GRAPH

Normalization is What Makes FHIR Binding Trivial

FAERS 2004–2026 Q2, de-duplicated by case ID

- 4,032,362 cases loaded into Neo4j
- 2,993,318 (74.2%) carry an outcome class code
- 11 node types · 9 relationship types · 8 indexes
- Worst-outcome severity ordering per case

The RxNorm spine

Drug →[HAS_INGREDIENT]→ Ingredient (RxCUI) →[IN_ATC_CLASS] → AtcClass4

478,424 drug-name strings collapse onto 9,381 ingredients and 715 ATC-4 classes.

WHY THIS MATTERS HERE

US Core requires RxNorm coding on Medication Request.

Our graph is RxNorm-normalized.

So a medication resource pulled from the EHR binds directly to Ingredient.rxcui — **no string matching, no mapping table.**

When a newly approved drug is missing from the graph:

A three-tier enrichment pipeline resolves it — 1) live RxNorm REST API; 2) a DailyMed-grounded local LLM; 3) an ungrounded tier that is flagged for human review and never auto-committed to KG.

HYBRID RETRIEVAL

Three Channels, Fused by Graph-aware Re-ranking

① LEXICAL

BM25 over case text

- Exact-term matching on drug and
- indication strings

- Catches precise nomenclature that
- dense embeddings blur together

② SEMANTIC

Fine-tuned BGE-M3 in Milvus

- 500K-triplet contrastive fine-tune
- of BAAI/bge-m3

- HNSW, ef = 256, top 200 by cosine

- 214× Top-1 enrichment over random

③ RELATIONAL

Graph-constrained expansion

- A case qualifies only if it shares an
- active ingredient AND an indication

- Matched on the normalized ingredient,
- so brand variants still resolve

- “Fan-out” bounded traversal

FUSED RE-RANK: All three channels merge and re-rank on a composite of normalized text relevance. The top-K subgraphs go to a locally hosted Qwen3.5-9B (LM Studio), one model for both the summary and natural language→Cypher query; its static prompt prefix is KV-cached, so a follow-up only computes the new question without prefix. Every generated statement traces to its source FAERS case IDs.

FHIR INTEGRATION

FHIR Supplies the Query Context

FHIR resource	Element	Query input
Patient	birthDate, gender	age, sex
Observation	LOINC 29463-7	body weight
MedicationRequest / MedicationStatement	medicationCodeableConcept (RxNorm)	**drug → Ingredient.rxcui**
Condition	code	indication

SMART on FHIR

EHR launch, OAuth 2.0 + PKCE, patient-context scopes. R4 (4.0.1), US Core 6.1.0, SMART App Launch 2.0.

Returned to the EHR

RiskAssessment / DetectedIssue summarizing the matched cohort, linked to supporting FAERS case IDs.

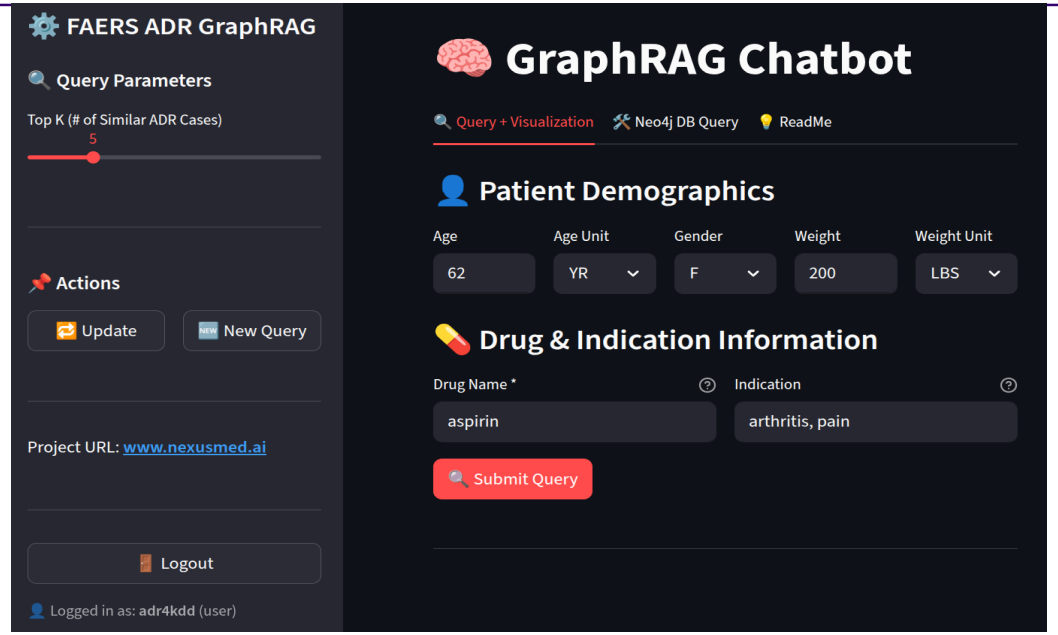
No PHI retained — resources are session-only.

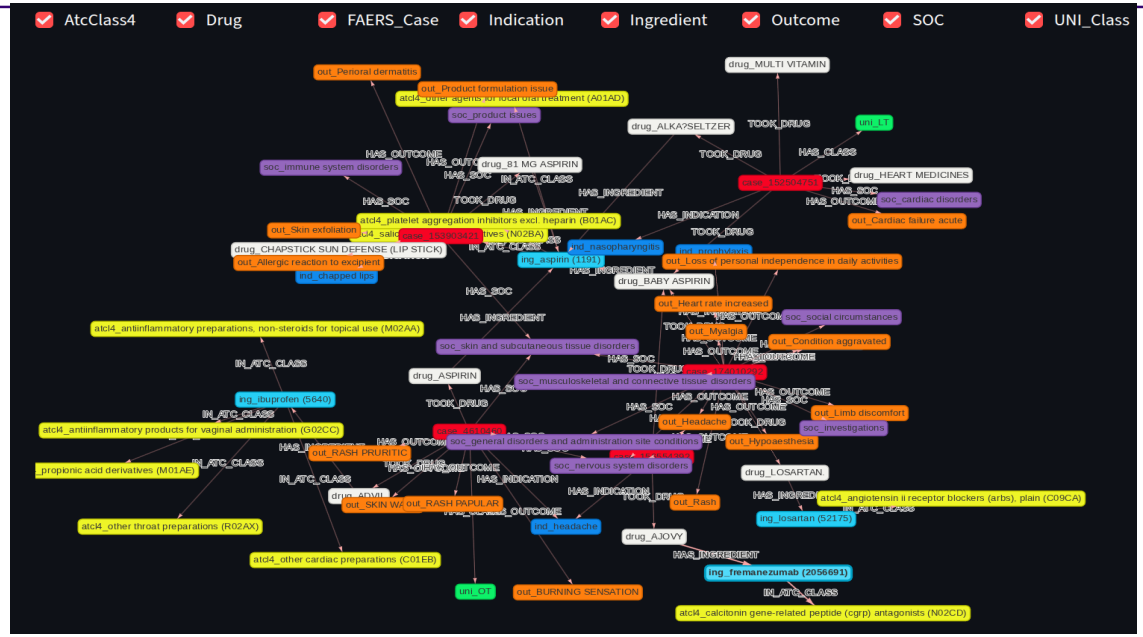
The corpus itself is not FHIR.

FAERS is public FDA data ingested offline into Neo4j and Milvus. FHIR carries the patient context in — it is the interface, not the evidence base.

THE APPLICATION

Every Answer Traces to Cases and PubMed Publications





What is on screen

- Ranked cohort of matched real-world cases, with a streamed local-LLM summary
- Interactive subgraph: which drug → indication → outcome paths drove the match
- Every assertion resolves to FAERS case IDs; each natural language query result links its top-3 PubMed papers
- Two mode tab: ask in natural language or Cypher statement; generated queries are schema-checked, self-corrected, and run read-only

Let Us Have a Quick DEMO



graphrag.nexusmed.ai

DISCLAIMER

What we don't claim

- Causation — FAERS has reporting bias
- Clinical validation — not yet evaluated in live workflow
- Outcome impact — no prescribing or patient effect measured

It is a hypothesis-generation and evidence-retrieval tool.

Huggingface Share : ADR Embedding Model

Selected Publications:

1) David Guo, Nishant Vishwamitra, Kim-Kwang Raymond Choo (2026). *Attention-Weighted Retrieval-Augmented Data Engineering for ADR Severity Classification*, EMNLP 2026 Industry Track Conference, ACL_EMNLP , Github Repo

2) David Guo, Kim-Kwang Raymond Choo (2025). *Applications of Federated Large Language Models for Adverse Drug Reaction Prediction: A Scoping Review*, Journal of Medical Internet Research, Doi: 10.2196/68291