


Enabling Reproducible AI-Driven Research via a Unified Intelligence Layer provides a ground truth that can be leveraged by AI and humans

Title: Separating Intelligence from Application: A Standardized Approach for AI-Driven Evidence Generation

Background: Research workflows require the use of concept sets, phenotypes and cohort definitions, which are often implemented in isolation of one another. To enable AI for more advanced workflows this context needs to be brought together to form an intelligence layer. This enables standardized reusable knowledge resources that are leveraged by Large Language Models (LLMs) and Generative AI that can be used to enhance user experience within a broad range of software applications from study inception to evidence generation. Imposing guardrails on how definitions are interpreted by the LLMs, limiting hallucinations and enhancing reproducibility.

Methods



Results

- Standardized AI Context Engineering: Intelligence Layer provides context (e.g., specific phenotype logic and transformation definitions), reducing hallucinations and improving the accuracy of AI-generated evidence.
- Logic Reusability: Definitions are authored once and exposed as services to multiple frontend applications, ensuring that an AI and a human researcher are always looking at the same version of a definition.
- Auditability: Every metric or cohort produced is linked directly back to a specific version of a definition within the Intelligence Layer.

Conclusions

- By implementing an intelligence layer in conjunction with the OMOP Common Data Model, the barriers to incorporating AI-driven workflows into applications are significantly simplified.
- Dual-layer approach provides standardized context to the AI, which limits hallucinations and the propagation of biases inherent in Real World Data, thereby improving the accuracy of evidence generation.
- This system ensures that all logic is reusable and every output is fully auditable

