# Federated Clinical Question-Answering

Alice Chen · Emily Jiang · Lalana Kagal

## Motivation

- Electronic Health Records (EHRs) are widely adopted in healthcare
- EHR text is unique
  - private, dynamic, distributed, jargon-dense
- LLMs pose their own challenges
  - hallucination, interpretability
- Goal: **Allow clinicians to ask natural language questions about trends in their patients' health records**

## RAG

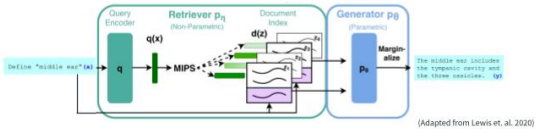


(Adapted from Lewis et. al. 2020)

- Retrieval-Augmented Generation: improve accuracy of LLM responses
- Generator answers query using docs from retriever
- RAG helps us use LLMs for the medical domain **without fine-tuning**!

## Architecture Diagram

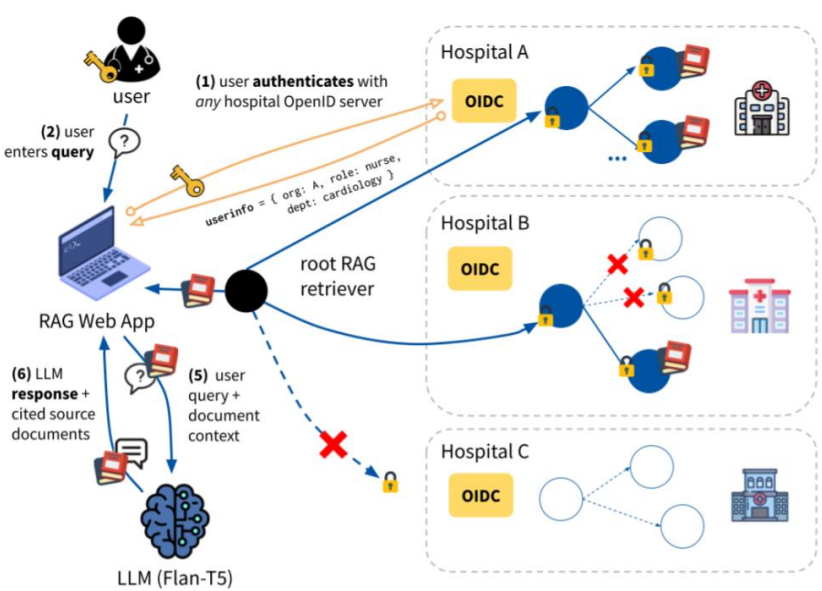


System architecture diagram for our federated + secure retrieval-augmented generation (RAG) over EHRs.

## Solution

- Extend RAG retrieval along the axes of **federation** and **security**
- Federation: query propagated across decentralized EHR database network
  - EHRs aggregated up hierarchy
- Security: Attribute-Based Access Control policies filter access at hospital, dept., EHR granularities
- Security: trusted OIDC server per hospital authenticates users

## Contributions

1. Use of RAG to synthesize trends over clinical data in distributed storage
2. Hierarchical retriever design with local document search at leaves and routing and aggregation at non-leaves
3. Access control mechanisms to uphold privacy guarantees of patient data
4. Creation of a clinical trend dataset for evaluating EHR QA systems