



The Pathways Project

Intelligent document search for pediatric clinical decision support



Kwaku Agyapong '26 · Ki Tae Hwang '26

CT Children's Medical Center × AIQUI Sandbox | Advisor: Prof. Kenneth Kousen | Trinity College Senior Seminar

Project Goal

Build a pathway-aware RAG chatbot that helps clinicians and staff search Connecticut Children's pediatric clinical pathways with fast, relevant, and source-grounded responses.

Problem and Use Case

Clinical pathways are dense, multi-document resources. The project focuses on turning guideline libraries into a searchable assistant that keeps answers tied to the selected pathway.

67

PATHWAYS

30

PAGES / PATHWAY

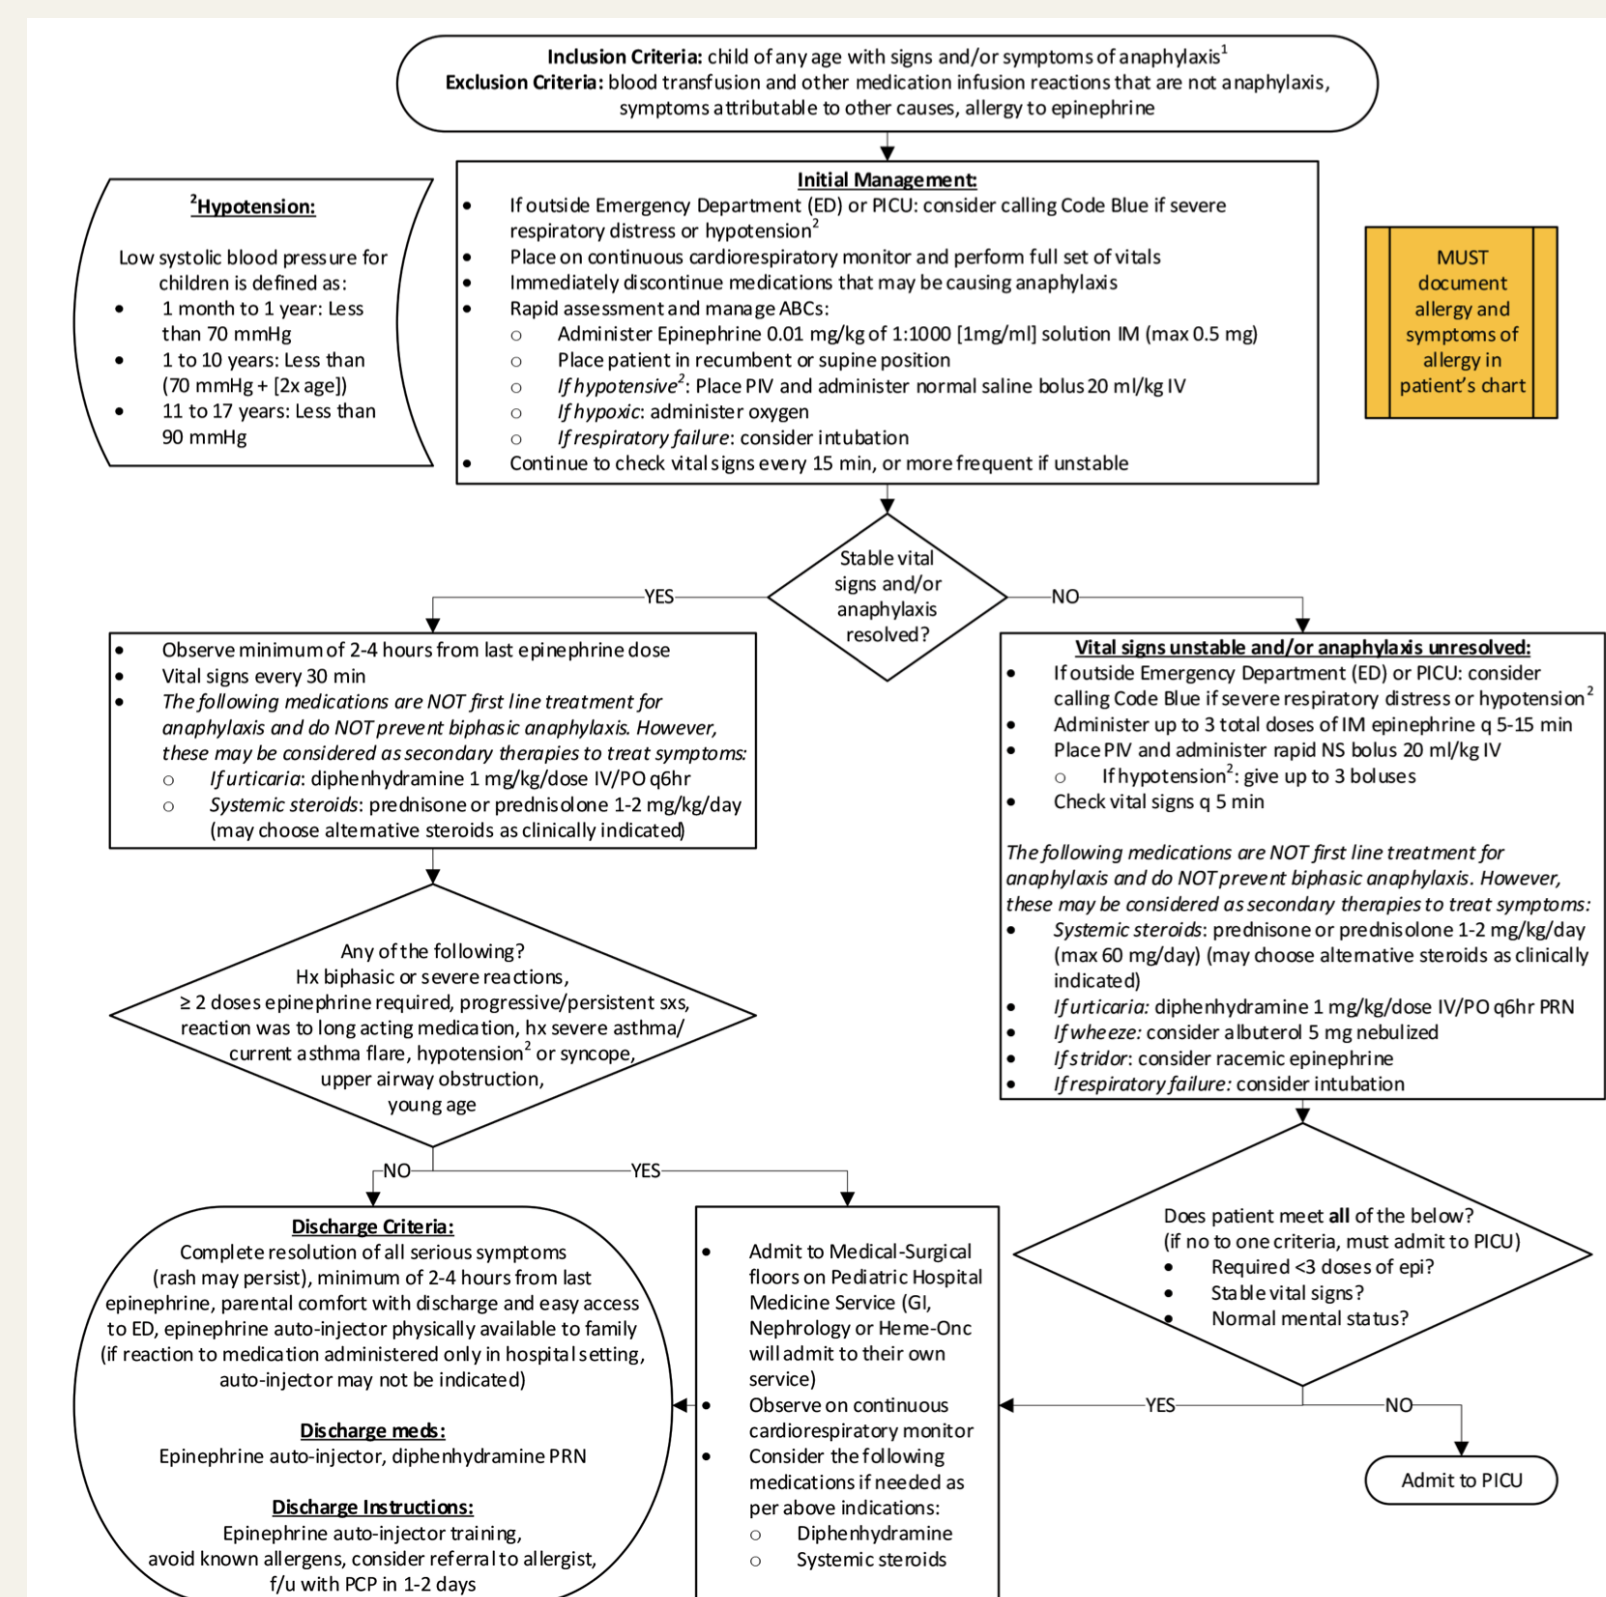
2010

SOURCE PAGES

Why this matters

- Clinical users need quick navigation across algorithms, educational modules, and related pathway PDFs.
- General-purpose LLM answers are not enough; retrieval must be limited to the selected clinical context.
- The system should reduce document-hunting while preserving clinical judgment and source visibility.

Clinical Pathway Source



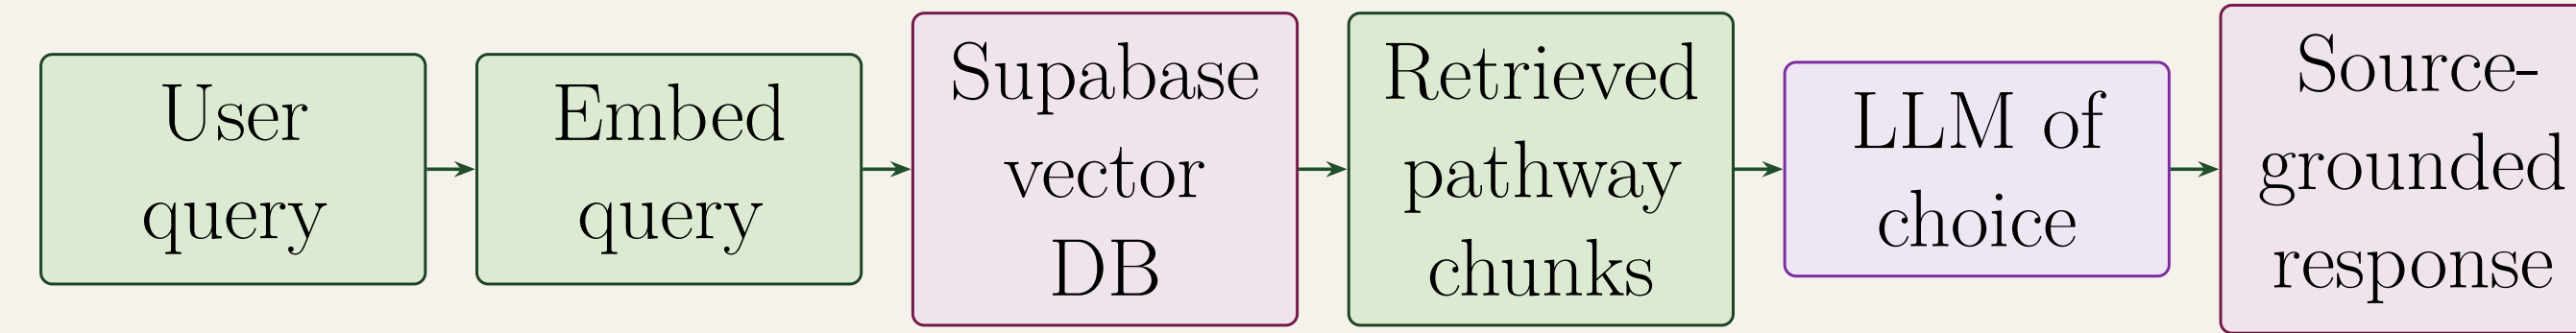
Anaphylaxis pathway: triage, treatment, and disposition algorithm.

Success Criteria

- Search is scoped to a chosen pathway, not the open web.
- Retrieval returns precise chunks with pathway and page metadata.
- Answers are presented alongside their source context.
- Monitoring and feedback enable iterative quality improvement.

Baseline RAG Pattern

Retrieval augmented generation turns a user query into a vector search over embedded documents, then gives the LLM the retrieved context before it answers.



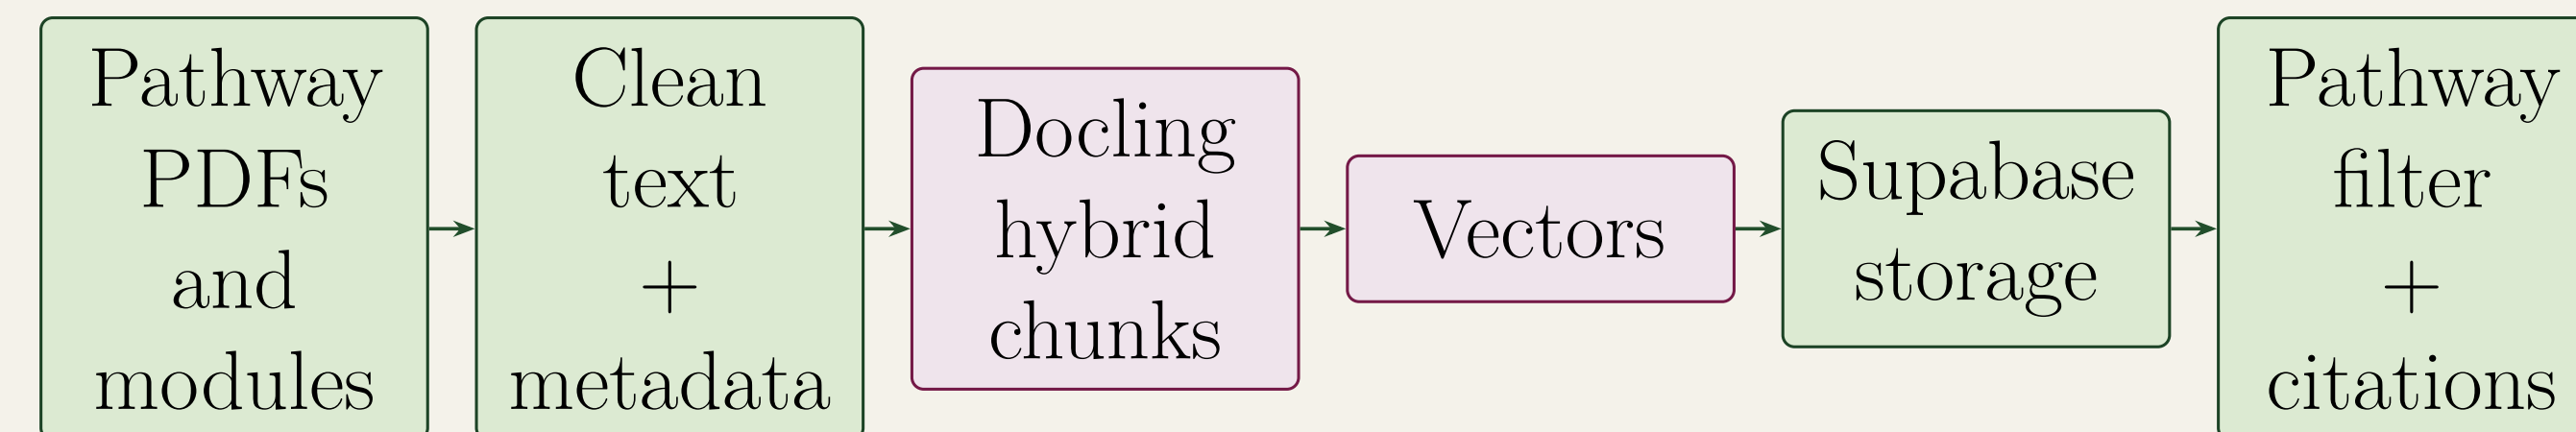
Query embedding → vector search → retrieved context → response generation.

Why not just a general-purpose LLM?

- A general LLM hallucinates specifics; clinical answers must be traceable to a pathway document.
- Open-web retrieval returns irrelevant or outdated guidance instead of the hospital's own pathway.
- Ground-truth sources at CT Children's live in versioned PDFs, not in the pre-training corpus.

Heavy-Digestion Pipeline

Raw PDFs carry broken tables, split paragraphs, and embedded figures that a direct ingest would dump into the index as noise. The heavy-digestion pipeline does the expensive document work *once*, up front, so retrieval at query-time is fast and precise.



Document transformation → chunking → embedding → vector storage → pathway filtering.

Design Principle

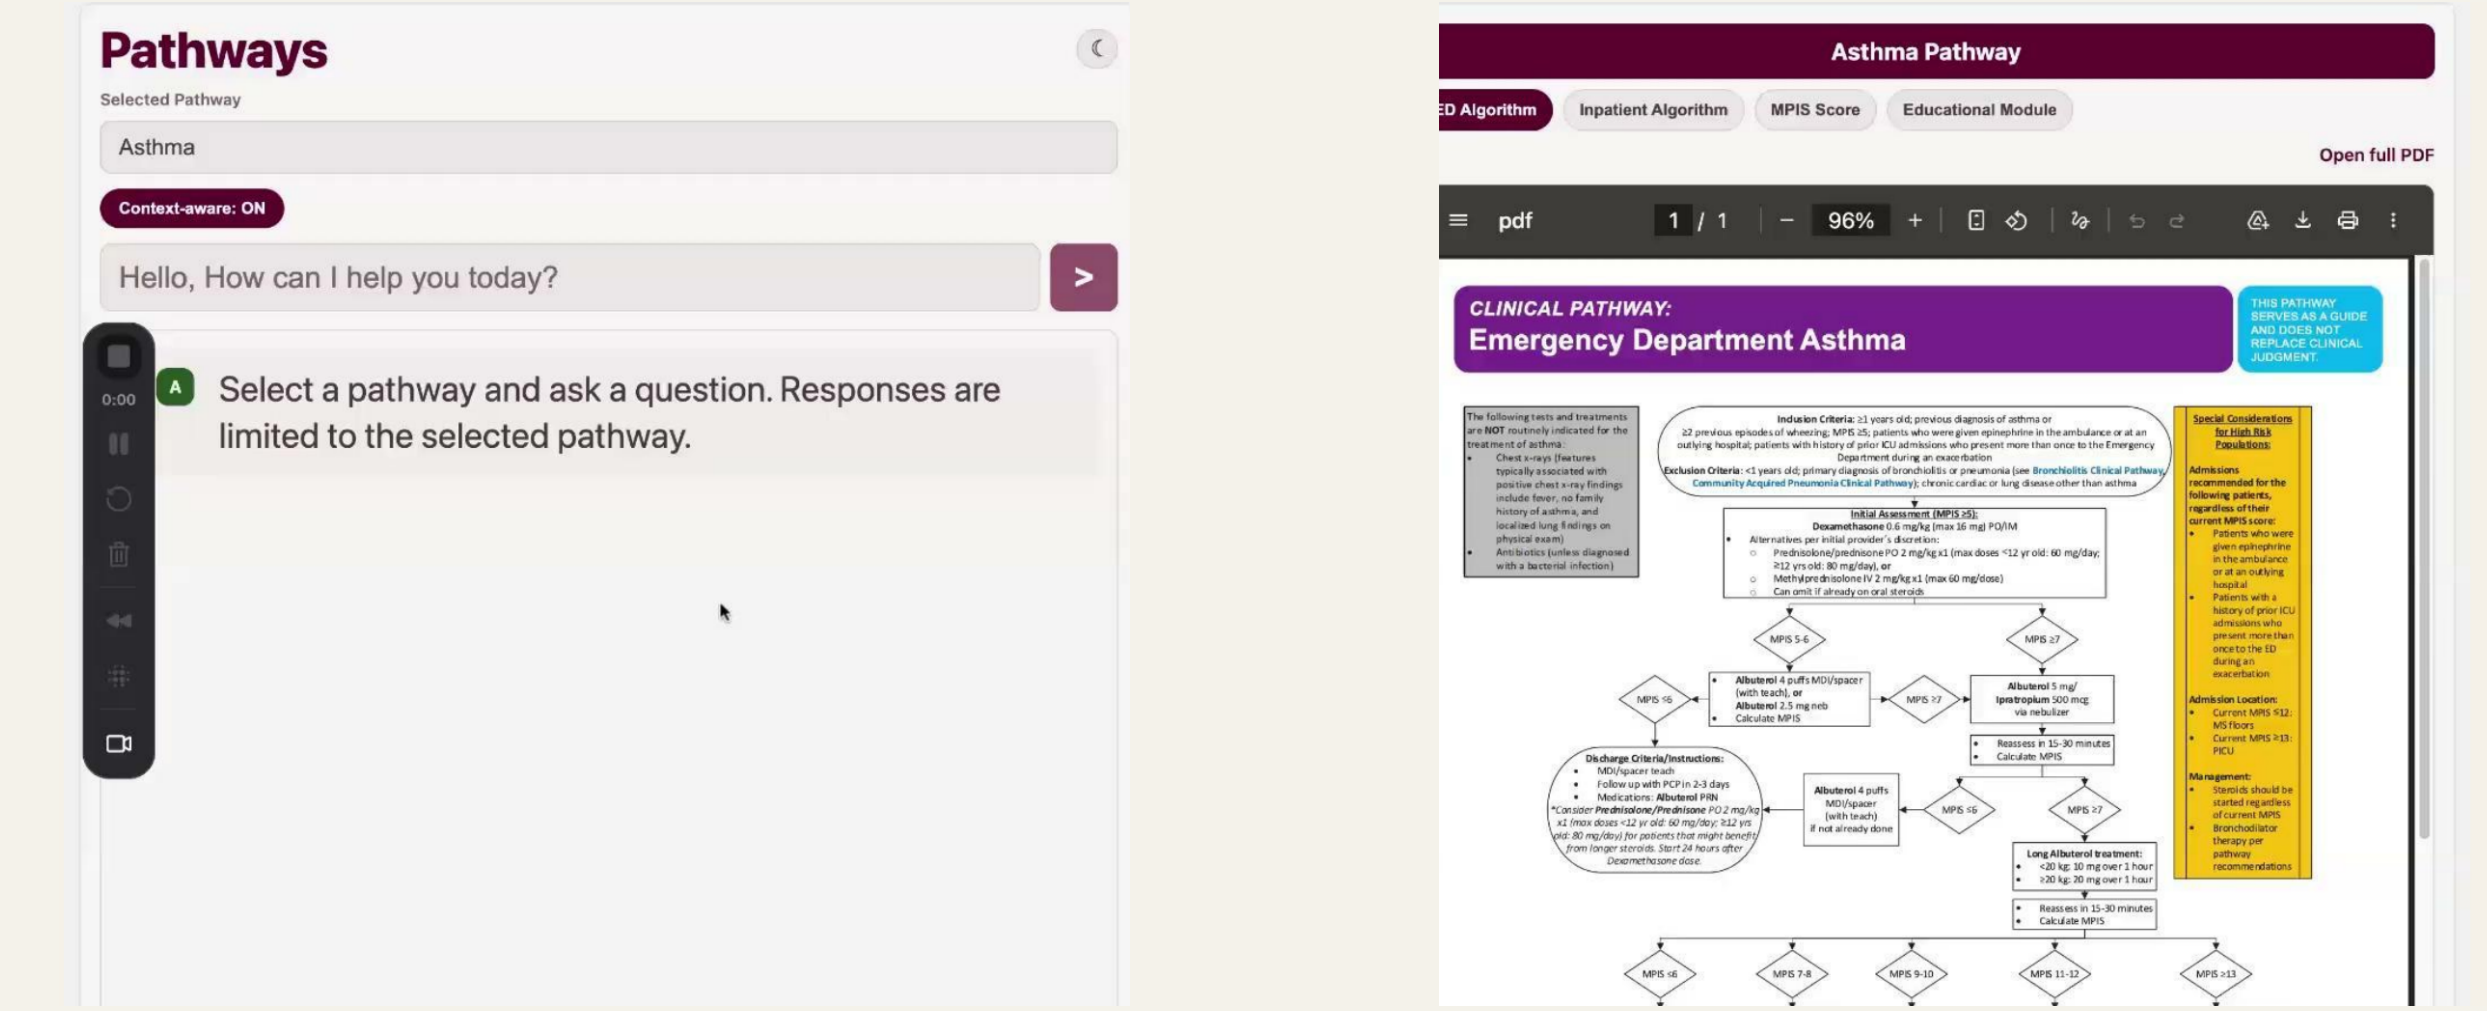
Do more document work *before* retrieval, so the chatbot can do less guessing *during* clinical question answering.

Direct Ingest vs. Heavy Digest

DIMENSION	Direct ingest	Heavy digest (ours)
Chunking	fixed char windows	document-aware (~512 tok)
Metadata	none / filename	pathway, section, page
Retrieval	whole corpus	pathway-filtered
Fidelity	lossy, generic	traceable, specific

Prototype Interface

The demo lets a user select a pathway, ask a question, and inspect the source PDF in one workflow.



Pathway-scoped chat, context-aware mode

Embedded PDF viewer, tabbed modules

- Pathway selection constrains retrieval to relevant source material.
- Context-aware mode supports natural follow-up questions.
- PDF tabs expose ED/inpatient algorithms, scores, and educational modules.

Answer Quality Across Rounds

Three evaluation rounds against a fixed 15-question clinical benchmark (Migraine, CAP, Asthma ED). Each round changed one layer of the stack—chunking, embedding, then pathway tagging—so gains could be attributed to specific decisions.

QUESTION	R1	R2	R3
CAP: antibiotic course length	5%	90%	92%
CAP: admission criteria	10%	35%	85%
Asthma: MPIS score calculation	58%	65%	80%

R1: MiniLM 384-dim, fixed chunks. R2: MedCPT 768-dim, semantic chunks + Claude image-to-text. R3: MedEmbed 1024-dim + pathway tagging.

- **R1→R2: the step change.** Semantic chunking and a clinical encoder drove the biggest gains—content trapped in figures was recovered by Claude image-to-text.
- **R2→R3: structure over dimensionality.** 1024-dim was marginal; **pathway tagging** unlocked CAP admission criteria by routing queries to the right document subset.
- **Remaining failures are semantic**, not structural—ambiguous clinical terms need query-side prompting.

Takeaways

- **Clinical alignment** — communication with CT Children's.
- **Prior art** — learning from Seattle and Arizona.
- **Impact** — building something that matters.

Next step: validate retrieval quality against pathway-specific questions, then connect monitoring and feedback data back into the document-processing pipeline.