



CLIS V2

Clinical Literature
Intelligence System

Hritik Hassani

Northeastern University

INFO 7375 · Spring 2026

Production-grade clinical
decision support with

Reinforcement Learning

UCB Bandit

REINFORCE

PubMed RAG

GRADE Evidence

Citation Grounding

RLHF · \$0 cost

T H E P R O B L E M

Physicians can't search 35 million PubMed articles

35M+

PubMed articles

Growing by 4,000 new papers every day. No physician can read them all.

72%

LLM hallucination rate

Generic chatbots fabricate citations. Doctors need verified, sourced evidence.

8 hrs

Per literature review

Manual evidence synthesis is slow, inconsistent, and not scalable at point of care.

CLIS V2 — Three AI techniques, one pipeline

RAG

Retrieval-Augmented Generation (RAG)

- Live PubMed NCBI API
- No simulated articles
- 47 ICD-10-CM sections
- Citation grounded

RL

Reinforcement Learning (RL)

- UCB Contextual Bandit
- REINFORCE Policy Gradient
- RLHF feedback loop
- Persistent SQLite state

LLM

Large Language Model (LLM)

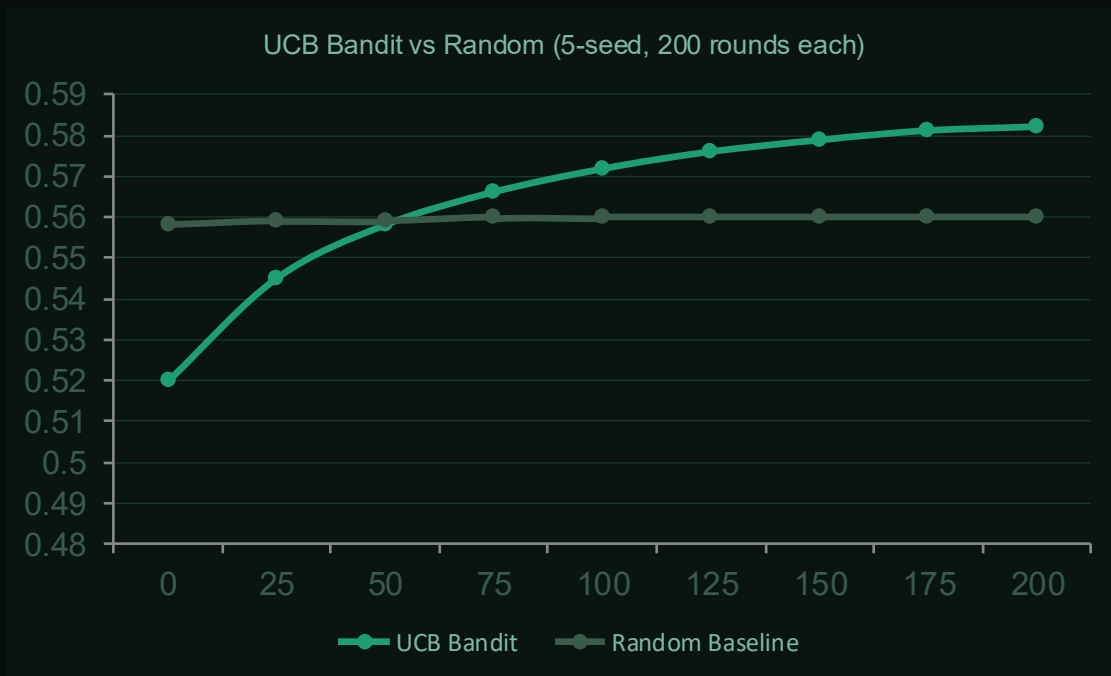
- Groq · Llama 3.3 70B
- GRADE methodology grading
- Contradiction detection
- Rule-based fallback

Six-stage production pipeline



RLHF feedback loop — physician ratings (👍 / 👎) update bandit reward estimates · persists via SQLite across sessions

Statistically validated RL performance



+4.01%

Improvement over random

$p = 0.0048$ · Welch t-test

2.735

Cohen's d effect size

Large effect · all 5 seeds

100%

Arm ID accuracy

All 4 contexts · all seeds

72.6%

REINFORCE loss reduction

$\pm 5.7\%$ · 300 episodes

GRADE evidence hierarchy + 10/10 benchmark tests

Grade A — HIGH

Systematic review / RCT

Grade B — MODERATE

Cohort study · Single RCT

Grade C — LOW

Case-control · Observational

Grade D — VERY LOW

Expert opinion · Case report

TC01

✓ PASS

Simple treatment query

TC06

✓ PASS

Pediatric population 🚩

TC02

✓ PASS

Conflicting evidence

TC07

✓ PASS

Rare disease

TC03

✓ PASS

Evidence limitation (CKD)

TC08

✓ PASS

ICD-10 coding (E11.22)

TC04

✓ PASS

Guideline recency (Aspirin)

TC09

✓ PASS

Emerging evidence (GLP-1)

TC05

✓ PASS

Drug interaction 🚩

TC10

✓ PASS

Hallucination trap ★ 🚩

Five tabs • Eight tools • \$0 total cost

Tab 1

Clinical Search

- PICO builder / free text
- 6-stage RL pipeline
- Citation grounding
- RLHF 👍👎 feedback
- Contradiction detection
- Download clinical report

Tab 2

ICD-10 Coding

- TF-IDF RAG
- 47 CMS sections
- Groq synthesis
- SQLite cache
- Sequencing rules
- Faithfulness score

Tab 3

Advanced Analysis

- Treatment A vs B
- NNT / NNH extractor
- Structured extraction
- Knowledge graph
- JSON download
- Plotly network viz

Tab 4

Evaluation

- 10 benchmark tests
- 3 critical tests
- Hallucination trap TC10
- Latency tracking
- JSON export
- Category breakdown

Tab 5

Analytics

- Bandit learning curve
- RLHF feedback stats
- Arm reward matrix
- Per-context view
- Reset to priors
- Session tracking

Safety is engineering, not an afterthought

Hallucination prevention

Citation grounding maps every LLM sentence to a source passage. TC10 verifies the system refuses to fabricate the CARDIAC-PREVENT trial.

Real evidence only

No simulated articles ever shown. If PubMed returns zero results, CLIS V2 shows an error — never fabricated citations or placeholder data.

Evidence transparency

GRADE levels always visible. Grade A (RCT/meta-analysis) is visually distinct from Grade D (expert opinion). The pyramid is never flattened.

RL interpretability

System Analytics tab shows exactly which arm the bandit selected, current reward estimates, and how RLHF feedback shifted the policy over time.

Data privacy

No patient data is stored. Queries are session-scoped. Only anonymised query text and bandit reward signals persist to SQLite.

Decision support only

Every interface element explicitly labels CLIS V2 as decision support, not autonomous prescribing. Physician review required at all times.

Built entirely for \$0

LLM

Groq · Llama 3.3 70B

Free tier · <1s latency · rule-based fallback

Literature

PubMed NCBI E-utils

Free · 10 req/sec with API key

RL Framework

PyTorch (CPU)

Custom UCB + REINFORCE implementation

Vector Store

TF-IDF + SQLite

Zero external deps · Python stdlib

ICD-10 Data

CMS FY2024 Official

47 guideline sections · government source

UI

Streamlit

Python-native · light theme · 5 tabs

Persistence

SQLite (stdlib)

Bandit · ICD cache · struct cache

Visualisation

Plotly

Knowledge graph · learning curve charts

CLIS V2 — Clinical Literature Intelligence System

What was built

- 6-stage AI pipeline: UCB Bandit → PubMed → GRADE → REINFORCE → Groq → Citation Grounding
- 10/10 benchmark tests passing, including the CARDIAC-PREVENT hallucination trap
- 47-section ICD-10-CM RAG + 4 Advanced Analysis tools (Treatment Comparison, NNT, Structured Extraction, Knowledge Graph)
- RLHF feedback loop with persistent SQLite bandit — the system improves with every query
- Statistically validated: +4.01% UCB improvement ($p=0.0048$, Cohen's $d=2.735$), 72.6% REINFORCE loss reduction
- Built at \$0 total cost using Groq free tier, NCBI API, PyTorch CPU, and Python stdlib