FHIR-Slice: Field-Level Degradation and Retry Planning for Clinical Decision Support under Partial Retrieval

Shivam Gupta
Independent Researcher, Delhi, India
shivam1720406@gmail.com | ORCID 0009-0009-4576-0138

Extended abstract for NBER Applications of Artificial Intelligence in Healthcare, 2-3 June 2027. Requested format: nonarchival poster.

Acknowledgment and disclosure. No external research funding was received for this study. The author is founder of Siloed, an AI software consultancy. No Siloed product is evaluated.

Research question and setting

A clinical decision-support application can produce ordinary empty lists or eligibility results after an incomplete FHIR search. This creates an integration question: which summary fields remain equal to their complete-retrieval values, and which searches should be retried when recovery has a limited budget? We study the original clinical logic of a pinned public AHRQ Pain Management Summary application. The contribution is an executable audit and controlled-degradation workflow grounded in established CQL data requirements and program slicing [1-3].

Method

FHIR-Slice associates each output field with the resource types read by its transitive Expression Logical Model dependencies. It releases a field only when every required retrieval scope has completed, otherwise returning a withheld envelope distinct from a clinical null, false result or empty list. A conditional argument establishes agreement with the complete-input evaluator when dependencies are conservative, completion receipts are truthful, the evaluator is deterministic and its non-query environment is fixed. Fault injection enumerates 255 nonempty resource-type failure masks for each of ten fixture representations; single-resource and prefix omissions provide additional tests.

Selected results from the finite software-test census; these are field evaluations and planning instances, not independent patients.

Quantity	Recorded result
Changed clinical-field evaluations under continue-on-error	17,856 / 40,800
Clinical fields released by the dependency gate	13,264
Disagreements among released clinical fields	0
Fields retained under single-type failures	984 / 1,280
Suboptimal immediate-gain plans, primary workload	390 / 2,048
Suboptimal immediate-gain plans, historical extension	124 / 3,616

Recovery and independent checks

A search can have no immediate benefit even though it belongs to a combination that restores a useful field. The planner therefore optimizes recovered fields over the dependency graph. In the primary workload, immediate-gain greedy repair is suboptimal in 390 of 2,048 instances, with gaps of up to six summary fields. A structural extension to twelve historical public quality-measure artifacts finds suboptimal greedy plans in 124 of 3,616 instances. Exhaustive planning, an independently implemented dependency analysis, runtime retrieval traces and integer programming check different parts of the implementation. Field counts are an explicit engineering objective; they are not unsupported estimates of clinical utility.

Relevance to healthcare AI infrastructure

The findings locate an assurance boundary between data retrieval and downstream decisions. An AI-facing interface could preserve the distinction between completed evidence, withheld outputs and ordinary clinical emptiness, while using dependency-aware recovery to preserve useful unaffected information. The measured intervention here is deterministic CQL execution. Integration with language models, changes in model quality and downstream clinical outcomes were not evaluated. This boundary makes the proposed poster relevant to health-data infrastructure and interoperability while keeping its empirical claims tied to the audited system.

Scope and reproducibility

The ten representations contain nine unique developer-authored fixture identities from one application family. Repeated failures form a test census and do not estimate operational incident rates. The twelve-artifact extension broadens dependency structure, but is not a second end-to-end clinical runtime. Agreement is with the supplied complete fixture, not with a patient-level ground truth; false completion receipts, concurrent chart changes or unsupported expressions can invalidate the assumptions. The artifact contains acquisition scripts, a dependency lockfile, protocol, compiler, output gate, fault generator, local HTTP transport test, independent verification, saved results and figure code. These support regeneration without patient-server access or model API calls.

Presentation status

This is a nonarchival poster proposal for completed computational work. The related full manuscript has been submitted to IEEE ICHI 2027; no acceptance or archival publication is claimed. The poster request does not seek a second proceedings publication.

Selected references

[1] HL7 International. FHIR R4, version 4.0.1 (2019). https://hl7.org/fhir/R4/

[2] AHRQ CDS Connect. Pain Management Summary SMART on FHIR application. Audited revision 664dbad5c82458bdbe8cc8caa16dd3119d43821c. https://github.com/AHRQ-CDS/AHRQ-CDS-Connect-PAIN-MANAGEMENT-SUMMARY

[3] Weiser, M. Program slicing. IEEE Transactions on Software Engineering SE-10(4), 352-357 (1984). doi:10.1109/TSE.1984.5010248.