Original Paper

An AI-Integrated Relational Curriculum Mapping Framework for Alignment and Transparency in Medical Education: Implementation Study

Robert M Badeau[1,2], PhD, MSc, PMI-CPMAI; Jennifer Laury[1], MA; Danielle Dickey[1,3], EdD

[1]Office of Academic Affairs, Naresh K. Vashisht College of Medicine, Texas A&M University, Bryan, TX, United States

[2]Department of Humanities in Medicine, Naresh K. Vashisht College of Medicine, Texas A&M University, Bryan, TX, United States

[3]Department of Medical Education, Naresh K. Vashisht College of Medicine, Texas A&M University, Bryan, TX, United States

Corresponding Author:
Robert M Badeau, PhD, MSc, PMI-CPMAI
Office of Academic Affairs
Naresh K. Vashisht College of Medicine
Texas A&M University
8447 John Sharp Pkwy
Bryan, TX, 77807
United States
Phone: 1 979 436 0621
Email: rmbadeau@tamu.edu

Abstract

Background: Medical school curricula are managed through fragmented, static tools that prevent faculty and curriculum committees from seeing a unified view of content scope, sequence, and licensing examination alignment. Governance relies on manual reporting that is neither timely nor consistent, and no widely adopted platform incorporates a generative AI agent that queries a relational curriculum database to produce tailored, on-demand reports.

Objective: This study aimed to design, implement, and evaluate a 4-phase, AI-integrated relational curriculum mapping framework providing faculty self-directed curriculum search; cross-course alignment reports classifying content as redundant, complementary, or a gap, with recommendations; and alignment with United States Medical Licensing Examination (USMLE) content and accreditation requirements.

Methods: We implemented the framework during academic year 2025-2026. It captures session learning objectives (SLOs) annually; maps them to program objectives, course objectives, USMLE codes, Association of American Medical Colleges keywords, and Bloom's levels using a large language model with expert verification; centralizes verified data in a relational database (Airtable) with faculty read-only search; and uses a second AI agent via an API to generate on-demand reports on scope, sequence, redundancy, gaps, and USMLE alignment across all 16 preclerkship courses and all clinical clerkship rotations.

Results: Course directors gained self-directed curriculum access through Airtable Extensions search. Cross-course vertical alignment reports identified redundant, complementary, and gap findings and generated course-level recommendations. An analysis of 2 preclerkship courses identified 50 matched SLO pairs across 3 USMLE pathology domains within minutes. A USMLE Step 1 alignment report across all 16 courses identified a curriculum-wide pathology coverage deficit (42% vs a target of 45%-55%) that was descriptively consistent with institutional performance data (not formally tested). An accreditation query identified all preclerkship and clerkship content on a specified topic and generated Data Collection Instrument-ready documentation within seconds. A Medical Education Program Objective coverage analysis across all 9 clinical clerkships, generated within minutes, found documented coverage for 24 of 25 (96%) program objectives and 1 objective with no matching SLO in any clerkship, a gap not previously tracked. In a blinded validation study, 3 independent course directors each coded a 30-SLO sample; agreement with the database (Fleiss κ, 3 raters plus database) was substantial at the exact-code level (κ=0.67, 95% CI 0.55-0.78) and almost perfect at the USMLE domain level (κ=0.94, 95% CI 0.88-1.00).

Conclusions: An AI-integrated relational curriculum database with a faculty-facing search interface can convert curriculum governance from a periodic, retrospective activity into a continuous, self-service one, producing alignment and gap analyses within minutes rather than weeks. Workload effects have not been formally measured. Extension to assessment items and prospective correlation with licensing examination outcomes are planned.

XSL•FO
RenderX

(JMIR Med Educ 2026;12:e107727) doi: 10.2196/107727

KEYWORDS

curriculum mapping; AI; education, medical; learning objectives; curriculum development; databases, relational; program evaluation; faculty

Introduction

Medical school curricula encompass thousands of session learning objectives (SLOs) distributed across preclerkship and clerkship phases; yet, the tools most institutions use to manage this content have not kept pace with the complexity they must represent. Most rely on flat file formats, spreadsheets, disconnected syllabi, and separate assessment records that cannot capture the relational architecture of a modern curriculum [1,2]. The consequence is "curriculum opacity": course directors and faculty cannot readily determine what is taught in adjacent courses, at what cognitive level, or whether it is assessed [1].

The practical costs are significant. When faculty lack cross-course visibility, unintended content redundancy accumulates; course directors reteach concepts already addressed elsewhere because they have no systematic way to know otherwise. Barde et al [3] documented that medical students and faculty identified redundant theoretical content in more than 30 of every 100 surveyed topics, with redundancy directly associated with perceptions of curriculum inefficiency and reduced student engagement. Without cross-course visibility, the distinction between introducing a concept and teaching it to mastery is operationally unenforceable. Youm et al [4] proposed a 6-degree model of curriculum integration encompassing content, assessment, faculty, and community dimensions, each of which depends on a system that makes the full scope and sequence of the curriculum visible to those responsible for designing it. Without that system, integration remains aspirational.

Curriculum committees face a parallel challenge: answering governance and accreditation questions about scope and coverage requires manual analyses that take weeks and may be outdated on delivery [1,2]. Existing commercial platforms, including Elentra, Entrada, One45, and VidaNovaVLE, have improved on spreadsheet approaches, but critical limitations persist: most require significant staff time to input and maintain data. Faculty remain passive consumers of typically uninterpretable staff-generated summary reports rather than active participants in a living curriculum system [2,5]. While computational approaches such as topic modeling have demonstrated the analytical potential of curriculum data, these have been implemented as one-time research exercises rather than persistent operational infrastructure [6]. Large language models are increasingly recognized as capable of natural-language classification and synthesis tasks at a scale not previously feasible; yet, their application in medical education has concentrated on student-facing tutoring and content generation rather than institutional governance infrastructure [7]. No widely adopted platform incorporates a generative AI agent capable of faculty and instructors querying the live database and producing tailored, on-demand reports in response to specific committee queries.

Previously, we described and validated a granular, hierarchical coding library derived from the USMLE (United States Medical Licensing Examination) Content Outline, together with a proof-of-concept application mapping SLOs and examination items to that library [8]. That earlier work established the granularity and machine-readability of the coding structure but did not address institutional-scale operationalization, faculty search access, or accreditation linkage. The present study builds on that infrastructure, extending it into a complete framework: a relational database populated using the coding library, a faculty-facing search interface, and an AI-agent reporting layer generating alignment and accreditation documentation across an entire preclerkship curriculum.

Methods

Setting and Development

The Naresh K. Vashisht College of Medicine Relational Curriculum Mapping Framework was developed and implemented during academic year (AY) 2025-2026, transitioning from annual, manually assembled Microsoft Excel–based curriculum snapshots to a centralized, AI-integrated relational architecture that connects learning objectives to program outcomes in a single queryable system. The framework spans 16 preclerkship courses and all clinical clerkship rotations. It is grounded in constructive alignment theory [9]. The framework comprises 4 sequential phases (Table 1 and Figure 1).

XSL•FO
RenderX

Table 1. Framework components, tools, inputs, outputs, and Liaison Committee on Medical Education (LCME) elements addressed by each phase[a].

Phase	Tool or platform	Input	Output	LCME element
1. Session capture	Excel intake form	Session titles and SLOs[b] from course directors	Version-controlled session data; prior-year database reference distributed to CDs[c]	6.1, 8.3
2. AI-agent mapping	Claude Sonnet 5 + human expert review	SLOs, USMLE[d] code library (18 domains, step 1 and step 2) [8], AAMC[e] keywords, Bloom's levels	AI-proposed tags verified by subject-matter experts before database entry	7.1, 8.3, 8.7
3. Database and faculty access	Airtable cloud relational database	Verified SLOs, USMLE codes, AAMC keywords, Bloom's levels, CLOs[f], MEPOs[g]	Faculty read-only searchable database; Airtable Extensions search for cross-course curriculum discovery	6.3, 8.3, 8.8
4. AI-agent API reporting	Claude API connected to Airtable live database via REST[h] API	Natural language queries from curriculum committee or mapping unit	USMLE alignment dashboards, redundancy and gap reports with actionable recommendations by SLO, CLO, MEPO, and USMLE domain	6.3, 8.3, 9.9

[a]Standards cited from LCME Functions and Structure of a Medical School, 2022 edition.

[b]SLO: session learning objective.

[c]CD: course director.

[d]USMLE: United States Medical Licensing Examination.

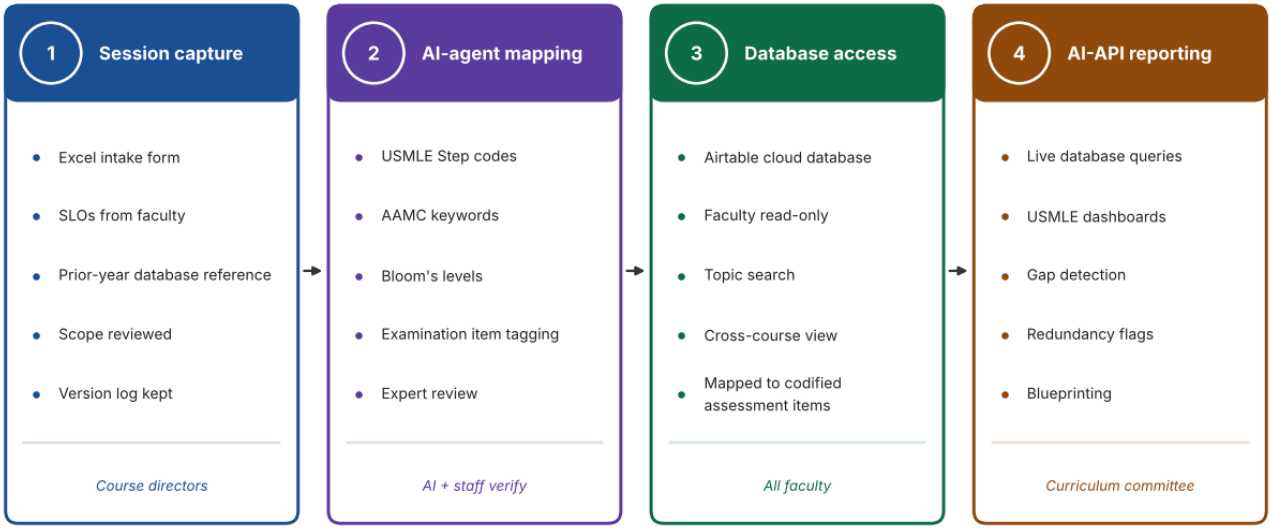
[e]AAMC: Association of American Medical Colleges.

[f]CLO: course learning objective.

[g]MEPO: Medical Education Program Objective.

[h]REST: Representational State Transfer.

Figure 1. Four-phase AI-integrated curriculum mapping workflow. Phase 1: annual session data capture via Microsoft Excel intake. Phase 2: AI-agent mapping of SLOs to USMLE STEP codes, AAMC keywords, and Bloom's Taxonomy levels with expert verification. Phase 3: centralization in the Airtable cloud relational database, with faculty read-only access, Extensions search for cross-course discovery, and preclerkship examination items mapped to codified assessment items linked to SLOs. Phase 4: AI-agent API reporting generating alignment dashboards, gap analyses, and redundancy reports by SLO, CLO, MEPO, and USMLE domain. AAMC: Association of American Medical Colleges; CLO: course learning objective; MEPO: Medical Education Program Objective; SLO: session learning objective; USMLE: United States Medical Licensing Examination.



Phase 1: Session Data Capture

Course directors submit SLOs annually via a standardized Excel intake form. Only SLOs are requested from faculty. Before finalizing SLOs for the incoming year, each course director receives read-only access to the prior AY's relational Airtable database, enabling intentional scope and sequence decisions grounded in direct visibility of what adjacent courses already teach, as well as review of the USMLE content coding, Association of American Medical Colleges (AAMC) keywords, and Bloom's Taxonomy levels assigned to their own course's SLOs from the prior cycle.

A well-functioning curriculum map also documents curricular changes systematically across AYs. Using an AI-API interface with Airtable, each annual submission cycle is version-controlled, with additions, deletions, and modifications to session sequence and SLO content logged against prior archived data. This longitudinal record supports Liaison Committee on Medical Education (LCME) Standard 8

XSL•FO
RenderX

expectations for curricular management: rather than reconstructing history from scattered records at a site visit, the curriculum committee has a continuously updated, queryable archive of how the curriculum has evolved.

Phase 2: AI-Agent Mapping

Each SLO is mapped to USMLE Content Outline codes, AAMC keywords, and Bloom's Taxonomy levels using structured, single-turn prompts submitted to Claude Sonnet 5 (Anthropic PBC; model identifier claude-sonnet-5, accessed May through September 2026) via its public API. Each prompt supplies one SLO together with the full USMLE coding library and a fixed set of classification rules and returns a single proposed code with a brief rationale; the workflow does not use retrieval-augmented generation or a multistep agentic pipeline at this stage, since the coding library is supplied directly in the prompt context rather than retrieved. Representative prompt templates, the classification rules used to select a code, and example inputs and outputs are provided in Multimedia Appendix 1. The framework is model-agnostic: any comparable large language model accessible via API can be substituted. All AI-generated tags are reviewed and verified by subject-matter experts [7].

The USMLE coding library and tagging methodology underlying this process were described and validated in the previous study reporting the construction and preclerkship application of the coding library [8].

AI-Tagging Validation

To quantify the accuracy of phase 2, a blinded validation study was conducted. A random sample of 30 SLOs was drawn across all 16 tagged courses, stripped of their AI-proposed and staff-verified USMLE codes, and distributed with only the session title and SLO text to 3 course directors with USMLE-content expertise who had not been involved in the original tagging of the sampled items. Each rater independently assigned a USMLE content code to every SLO using the same coding library, without sight of the database's existing code or of the other raters' responses. Rater-assigned codes were compared against the corresponding database codes and against one another, classified at 2 levels of granularity: an exact match (identical code) and a domain-level match (same top-level USMLE domain but a different specific subcode). Items a rater could not confidently code from the text alone were excluded from that rater's scored comparisons rather than treated as a mismatch. Agreement across the 3 raters and the database was summarized with Fleiss κ, computed separately at the exact-code and domain levels; 95% CIs were obtained by bootstrap resampling of the 30 items (3000 resamples). Pairwise Cohen κ between each pair of raters and between each rater and the database was computed as a secondary, descriptive comparison.

Phase 3: Relational Database and Faculty Access

Subject-matter expert–verified data are uploaded into Airtable, a cloud-based relational database. The framework is not limited to Airtable: any cloud-based relational database with API access can serve as the infrastructure backbone.

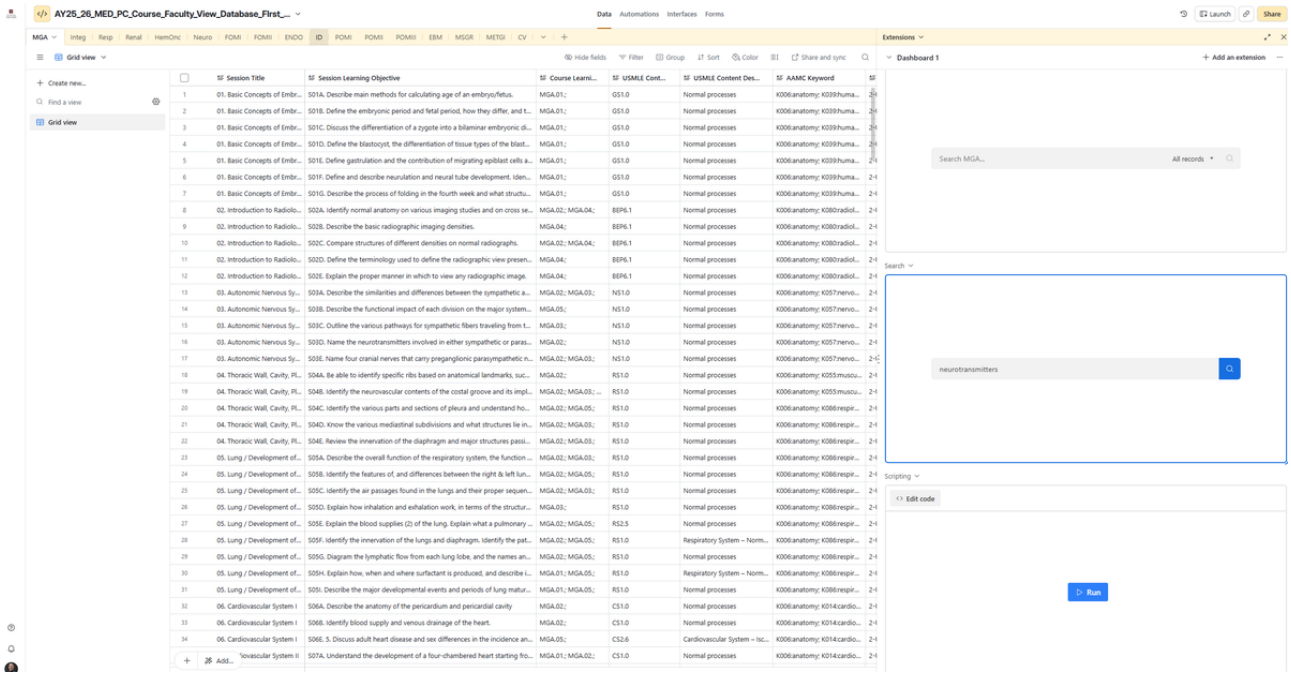
A read-only shared link is distributed to all course directors and faculty (Figures 2 and 3). The interface is deliberately simple: all 16 preclerkship courses are laid out as tabs across the top of a single screen, visible to any user immediately upon opening the database, without submenus, additional logins, or navigation beyond a single click. The key feature enabling cross-course discovery is Airtable's Extensions search, which allows any faculty member to query the entire curriculum database using free-text search and filters without technical training. A course director types a clinical concept and sees, within seconds, every SLO, CLO, Medical Education Program Objective (MEPO), and Bloom's Taxonomy level across all 16 courses addressing that concept.

The preclerkship database is mirrored by a clerkship database aligned to the clerkship core objectives, the Accreditation Council for Graduate Medical Education core competencies, clerkship rotation–specific learning objectives, institutional program objectives (MEPOs), and the USMLE Step 2 Clinical Knowledge framework (Multimedia Appendix 2). This cross-phase architecture operationalizes vertical curriculum integration [10].

The database contains curriculum content only: session titles, learning objectives, coding tags, and course-level metadata submitted by course directors. It contains no student records, grades, enrollment data, or other information covered by the Family Educational Rights and Privacy Act (FERPA), and access is accordingly governed by institutional data-use policy rather than FERPA-specific controls. Access is restricted to a read-only link distributed to faculty and curriculum staff through the institution's existing single-sign-on authentication; no public or student-facing access exists at this stage. Phase 2 tagging submits only deidentified SLO text, containing no patient or student information, to the external AI vendor's API; the vendor's standard data-retention and training-opt-out terms for API traffic apply, consistent with the institution's existing data-processing agreements for other cloud services.

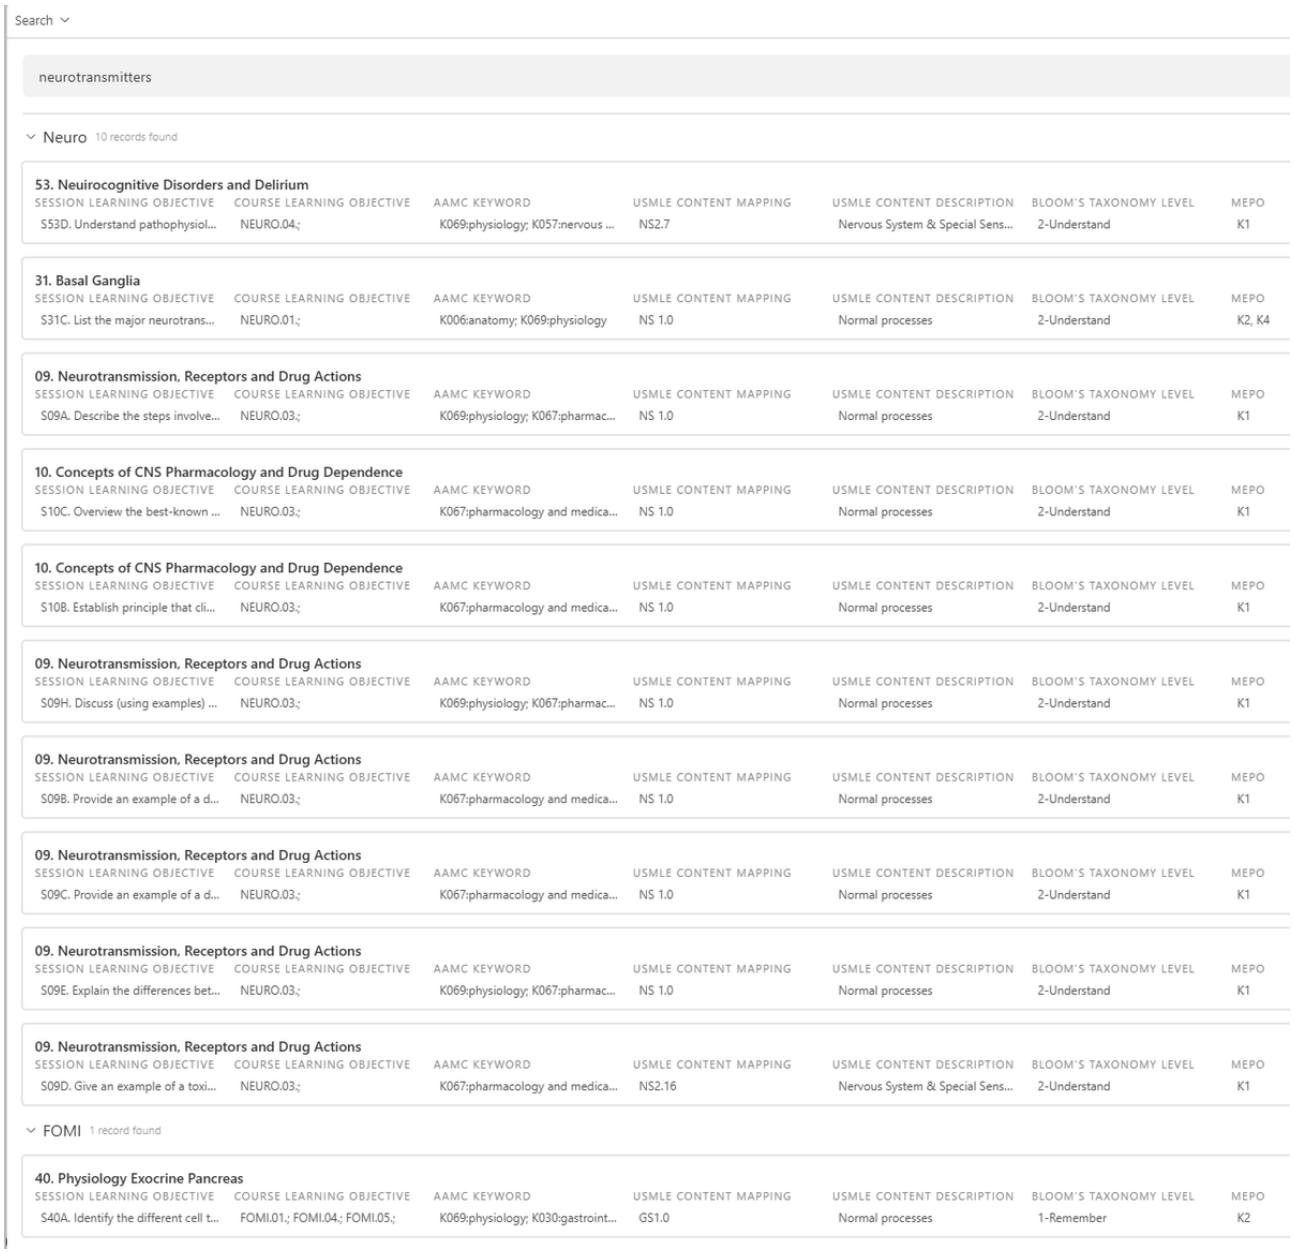
XSL•FO
RenderX

Figure 2. Screenshot of the faculty-facing read-only Airtable preclerkship curriculum database layout. All 16 preclerkship courses appear as tabs across the top of a single screen, immediately visible without scrolling or navigating a menu; the course view beneath displays session titles, session learning objectives, and mapped USMLE codes, AAMC keywords, and Bloom's Taxonomy levels for the selected course. AAMC: Association of American Medical Colleges; USMLE: United States Medical Licensing Examination.



XSL•FO
RenderX

Figure 3. Screenshot of the faculty-facing read-only Airtable preclerkship curriculum database. The result of an Extensions search for the term "neurotransmitters," returning every matching session learning objective across all 16 courses, organized by course, with the corresponding USMLE code, AAMC keyword, and Bloom's Taxonomy level displayed for each match. Faculty search across all 16 preclerkship courses by session topic, USMLE code, AAMC keyword, or Bloom's Taxonomy level without staff mediation. AAMC: Association of American Medical Colleges; USMLE: United States Medical Licensing Examination.



Phase 4: AI-Agent API Reporting

An AI agent interfaces directly with the live database via the Airtable Representational State Transfer API, accepting natural language queries from curriculum staff or the committee. This stage is a retrieval-based workflow rather than a multistep autonomous agentic pipeline: the query retrieves the relevant matched records directly from the relational database, and the retrieved records, not a separate search or planning step, are passed to Claude Sonnet 5 to be classified and written into a report. Report generation takes minutes; the complete pipeline including human expert review and committee-ready delivery is completed within minutes. No institution-level AI infrastructure is required.

A central phase 4 output is the longitudinal alignment report: a query comparing session titles and SLO text across a multicourse sequence that classifies each matched pair as redundant, complementary, or a gap, and generates a specific recommendation for each finding, such as consolidating redundant sessions, adopting cross-course numbering or cross-referencing conventions, or designating a single owning course for a given concept. Each report closes with an explicit caution that classifications involve interpretation and should be verified by course directors before curricular changes are made.

USMLE Step 1 Alignment Analysis

For each of the 16 preclerkship courses, discipline coverage percentages were calculated as the count of that course's SLOs coded to a pathology, physiology, or pharmacology USMLE

XSL•FO
RenderX

domain, divided by the total number of SLOs in that course. An SLO carrying more than 1 code was counted once toward each distinct discipline its codes mapped to, so a course's 3 discipline percentages do not necessarily sum to 100%; SLOs coded to a domain outside these 3 disciplines contributed to the denominator but not to any of the 3 numerators. Categorical status labels (on, below, above, and exceeds) were assigned by comparing each course's discipline percentage to the corresponding USMLE Step 1 Content Outline target range: below or above indicates the percentage fell outside the target range by less than 10 percentage points, and low gap or exceeds indicates a deviation of 10 percentage points or more. The composite alignment score (0-100) was calculated as 100 minus the sum, across the 3 disciplines, of each discipline's absolute percentage-point deviation from the nearest edge of its target range, floored at 0; a course exactly within all 3 target ranges receives a score of 100, and larger or more numerous deviations reduce the score proportionally. Evidence-Based Medicine and Medical Student Grand Rounds were scored against only the discipline targets applicable to their content, given their methodological and integrative focus rather than organ-system or disease content.

Ethical Considerations

Curriculum mapping and database development activities described in this manuscript constitute quality improvement and do not meet the regulatory definition of human subjects research [11].

Results

During AY 2025-2026, the framework was fully operational across all 16 preclerkship courses and all clinical clerkship rotations.

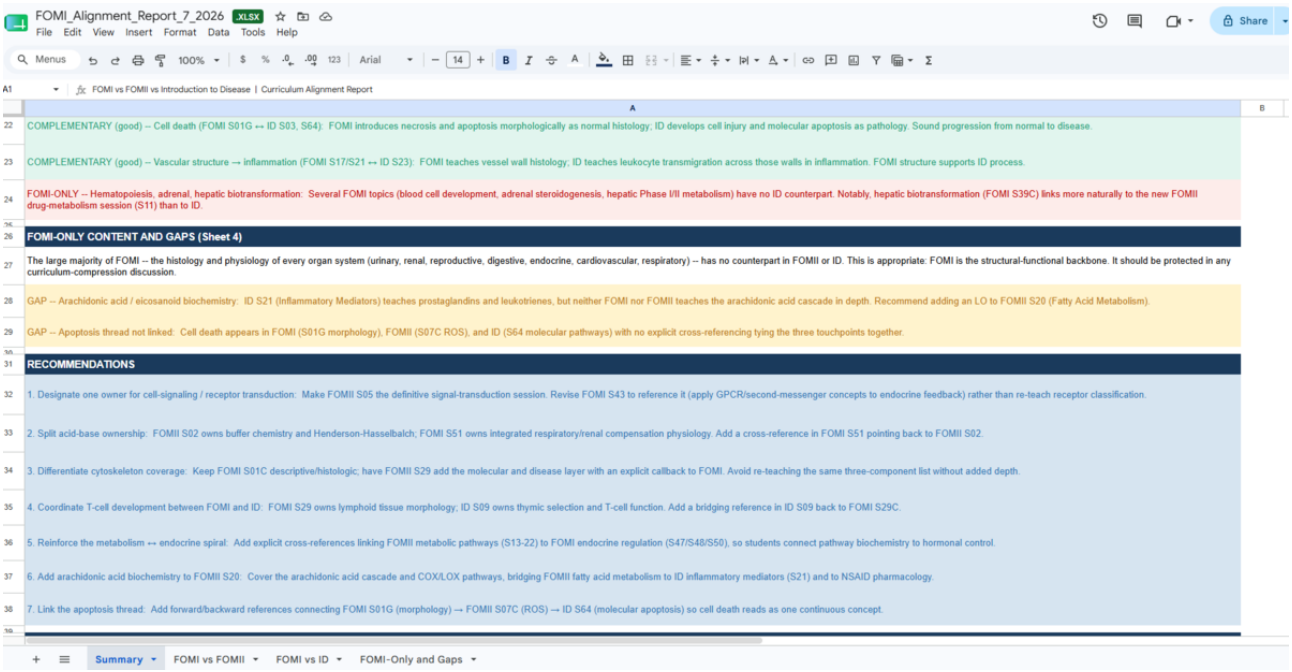
Faculty Access and Preventing Unintended Content Repetition

Read-only database access was distributed to all 16 preclerkship course directors and all clerkship discipline directors. Faculty used the Extensions search to verify topic coverage before finalizing annual SLOs, replacing informal colleague communication as the primary mechanism for cross-course coordination. Faculty also received AI-generated alignment reports classifying their SLOs against related courses as complementary, redundant, or a gap, with recommendations for revising session content to improve horizontal alignment; an example of one such report is shown in Figure 4.

The practical impact of cross-course visibility was demonstrated when the Curriculum Mapping Unit compared Foundations of Medicine I (FOMI, semester 1, AY 2025-2026) and Introduction to Disease (ID, semester 2, AY 2025-2026). A structured AI query identified 50 matched SLO pairs across 14 FOMI and 15 ID sessions spanning 3 USMLE pathology domains: cell injury and apoptosis, inflammation and repair, and neoplasia. Delivered within minutes, the report showed for the first time that FOMI's histology and cell biology sessions systematically introduce the structural substrate ID later formalizes as pathology; the FOMI cell organelles session alone connects to 10 downstream ID pathology SLOs. Course directors can now scaffold this progression explicitly rather than leaving vertical connections invisible to students, one practical response to the reteaching problem that curriculum opacity can create.

XSL•FO
RenderX

Figure 4. Screenshot of a longitudinal alignment report comparing FOMI, the revised FOMII, and Introduction to Disease, exported from the framework's database as a structured spreadsheet. Each row is labeled with 1 of 4 classifications: COMPLEMENTARY, indicating sound scaffolding from normal structure to disease process, illustrated here by FOMI's morphological introduction of cell death (necrosis and apoptosis) preceding Introduction to Disease's treatment of molecular apoptosis as pathology; GAP, identifying content taught in neither course, such as arachidonic acid and eicosanoid biochemistry; FOMI-ONLY, denoting foundational content, such as organ-system histology and physiology, that appropriately has no downstream counterpart and should be protected from curriculum compression; and RECOMMENDATIONS, a dedicated set of rows translating each finding into a specific action, such as designating FOMII as the definitive session for cell-signaling and receptor transduction with FOMI referencing rather than reteaching it or adding a learning objective to an existing session to close an identified gap. Each row includes the specific session codes involved and a brief rationale, providing course directors with an audit trail from finding to recommended action. Rows are color-coded by finding type. Green indicates a COMPLEMENTARY finding (sound scaffolding between courses), red indicates a GAP or FOMI-ONLY finding (content taught in neither course, or foundational content appropriately unique to FOMI), dark blue denotes section headers separating report categories, and unshaded rows under the RECOMMENDATIONS header list denote the specific corrective action for each numbered finding above it. FOMI: Foundations of Medicine I; FOMII: Foundations of Medicine II.



Alignment Reports With Actionable Recommendations for Course Directors

A parallel longitudinal alignment report compared the 3-course Practice of Medicine sequence (Practice of Medicine I [POMI], Practice of Medicine II [POMII], and Practice of Medicine III), finding that patient-centered interviewing was taught in overlapping form at least 5 times across POMI and POMII, and that POMI and POMII independently assign the same session number to 3 unrelated sessions, creating scheduling and assessment-tracking ambiguity distinct from content overlap. For each finding, the report generated a specific recommendation rather than a bare classification: consolidating redundant sessions, adopting a unified cross-course numbering scheme, revising physical examination objectives toward clinical reasoning rather than repeating skill-acquisition language, and adding sessions on telehealth and addiction medicine to address gaps identified through AAMC keyword tags with no corresponding taught session.

The FOMI-to-ID report similarly moved beyond classification to recommendation (Figure 4). For redundant content, such as cell-signaling receptor transduction and acid-base physiology taught in both FOMI and Foundations of Medicine II (FOMII), the report proposed designating a single owning course for each concept, with the nonowning course referencing rather than reteaching it, and for acid-base physiology, it proposed splitting ownership so that FOMII retains buffer chemistry and Henderson-Hasselbalch calculations, while FOMI develops integrated respiratory and renal compensation. For a gap in arachidonic acid and eicosanoid biochemistry, a prerequisite for inflammatory mediator content taught later in ID, it proposed a specific session and learning objective addition. A second gap, the apoptosis thread, appeared disconnected across all 3 courses: FOMI introduces its morphology, the revised FOMII covers its molecular mechanism, and ID addresses its pathological consequences, but no session in any course explicitly cross-references the other 2, and the report recommended adding forward and backward references so the 3 treatments read as one continuous concept rather than 3 unrelated topics. Course directors received the underlying matched-objective data alongside each recommendation and were advised to verify findings against current materials before making curricular changes.

USMLE Step 1 Alignment Analysis

A full USMLE Step 1 content alignment report spanning all 16 preclerkship courses was generated through a single natural language query and delivered to the curriculum committee within minutes (Table 2).

XSL•FO
RenderX

Table 2. Academic year 2025-2026 United States Medical Licensing Examination (USMLE) Step 1 discipline coverage and composite alignment scores across all 16 preclerkship courses[a].

Code	Course	Pathology (%)	Physiology (%)	Pharmacology (%)	USMLE status			Score out of 100
					Pathology	Physiology	Pharmacology	
POMII	POM II	44	30	18	On	On	On	99
METGI	Metabolism/Gastrointestinal	46	41	16	On	Above	On	99
RENAL	Renal	55	28	19	On	Below	On	97
POMIII	POM III	43	29	16	Below	Below	On	96
POMI	POM I	40	36	14	Below	On	On	93
ENDO	Endocrinology/Reproductive	51	31	28	On	On	Above	89
NEURO	Neuroscience	58	26	21	Above	Below	Above	89
FOMII	Foundations of Medicine II	38	42	12	Below	Above	On	88
RESP	Respiratory	49	52	12	On	Exceeds	On	84
MGA	Medical Gross Anatomy	37	44	6	Below	Above	Below	79
HEMEONC	Hematology/Oncology	62	22	22	Exceeds	Below	Above	77
FOMI	Foundations of Medicine I	33	48	10	Low	Above	On	73
ID	Introduction to Disease	54	14	26	On	Low gap	Above	71
INTEG	Integumentary System	68	18	10	Exceeds	Low gap	On	67
MSGR	Medical Student Grand Rounds	23	19	15	Low	Low	On	56
EBM	Evidence-Based Medicine	11	16	8	Low	Low	Below	45
AVG	All 16 courses	42	33	14	Below	On	On	—[b]
Target	USMLE Step 1 Blueprint	45-55	30-40	10-20	Reference ranges	Reference ranges	Reference ranges	—

[a]Generated via AI-agent API query against the Airtable curriculum database, May 2026. Composite alignment score (0-100) reflects simultaneous proximity to all 3 discipline targets. EBM and MSGR scored against applicable USMLE competencies only, given their methodological and integrative nature.

[b]Not applicable.

The report identified a curriculum-wide pathology coverage average of 42%, 3 percentage points below the USMLE Step 1 target of 45%-55%, with 6 of 16 courses below 40%. Physiology (33%) and pharmacology (14%) were both within target ranges. These findings were descriptively compared with institutional USMLE Step 1 performance data (National Board of Medical Examiners Step 1 Performance Profile, 2023-2025 cohorts) (unpublished institutional data, NBME Step 1 Performance Profile, 2023-2025 cohorts): a total of 34 of every 100 examinees scored below the national mean in pathology, compared with 33 of every 100 in physiology and 29 of every 100 in pharmacology. This pattern was only partly consistent with the coverage gap: pathology's below-target coverage coincided with the largest performance gap, but physiology showed a comparably large performance gap despite coverage within its target range, indicating that coverage alignment alone does not fully account for relative discipline performance. No formal statistical test of any of these associations was conducted. The report delivered prioritized, course-level recommendations alongside a composite alignment scorecard.

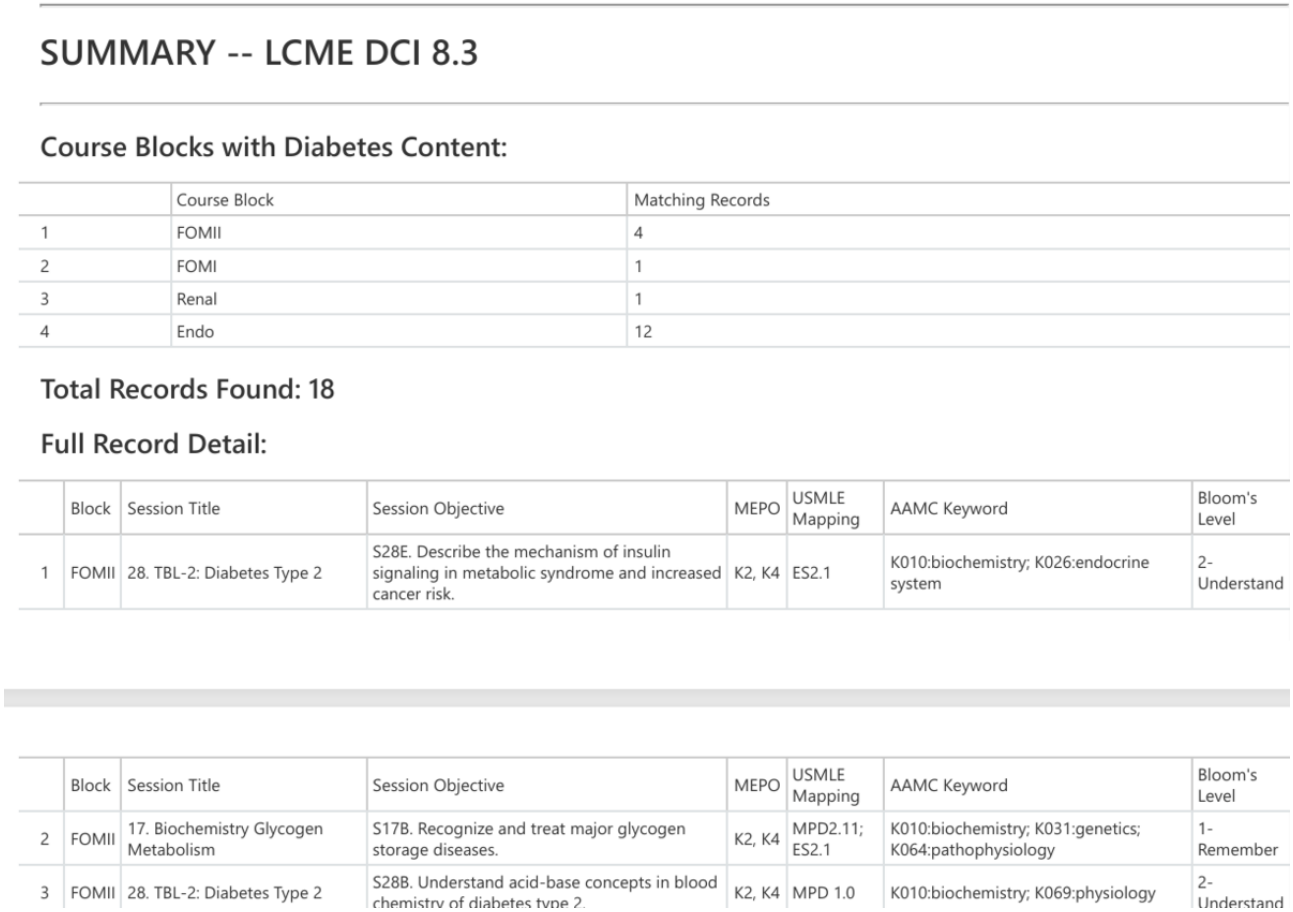
Accreditation Support

The relational database explicitly links every SLO to its parent course learning objective (CLO) and program objectives, addressing LCME Element 8.3, which requires faculty to define competencies and learning objectives for each curricular segment and link them to program objectives [12]. The read-only faculty interface evidences the LCME expectation that faculty can navigate and use the curriculum map. For Element 8.3, within seconds, staff generated a report, using a scripting extension in Airtable, to determine where a specified topic (diabetes) appeared across preclerkship and clerkship content for Data Collection Instrument reporting (Figure 5). The database supports accreditation readiness and gives the curriculum committee a tool for strategic curriculum planning: data-informed decisions about scope, sequence, and content alignment that previously required weeks of manual analysis by staff are now available rapidly and on demand, enabling the committee to identify gaps proactively. The framework was developed and is sustained without external funding or a dedicated commercial curriculum-mapping platform, relying instead on general-purpose, subscription-based tools (Airtable

XSL•FO
RenderX

and the Claude API) that are themselves proprietary but require no institution-specific licensing or custom development.

Figure 5. Example of an AI agent–generated LCME Standard 8.3 documentation report showing all preclerkship and clerkship course blocks containing diabetes-related content, including matching course blocks, record counts, and session-level detail with mapped MEPOs, USMLE codes, and AAMC keywords. Generated in seconds using an Airtable scripting extension from a single free-text query. AAMC: Association of American Medical Colleges; LCME: Liaison Committee on Medical Education; MEPO: Medical Education Program Objective; USMLE: United States Medical Licensing Examination.



SUMMARY -- LCME DCI 8.3

Course Blocks with Diabetes Content:

	Course Block	Matching Records
1	FOMII	4
2	FOMI	1
3	Renal	1
4	Endo	12

Total Records Found: 18

Full Record Detail:

	Block	Session Title	Session Objective	MEPO	USMLE Mapping	AAMC Keyword	Bloom's Level
1	FOMII	28. TBL-2: Diabetes Type 2	S28E. Describe the mechanism of insulin signaling in metabolic syndrome and increased cancer risk.	K2, K4	ES2.1	K010:biochemistry; K026:endocrine system	2-Understand
2	FOMII	17. Biochemistry Glycogen Metabolism	S17B. Recognize and treat major glycogen storage diseases.	K2, K4	MPD2.11; ES2.1	K010:biochemistry; K031:genetics; K064:pathophysiology	1-Remember
3	FOMII	28. TBL-2: Diabetes Type 2	S28B. Understand acid-base concepts in blood chemistry of diabetes type 2	K2, K4	MPD 1.0	K010:biochemistry; K069:physiology	2-Understand

Clerkship MEPO Coverage Analysis

To evaluate whether the framework's diagnostic capability extends to the clerkship phase, an AI-agent query generated a full coverage report tracing all 25 institutional MEPOs against SLOs across all 9 clinical clerkships (emergency medicine, internal medicine, family medicine, pediatrics, psychiatry, radiology, neurology, obstetrics and gynecology, and surgery). The report was generated within minutes from the relational clerkship curriculum database (Table 3).

Of the 25 MEPOs evaluated, 24 (96%) had documented SLO coverage in at least 1 clerkship, and 20 (80%) were covered across 5 or more of the 9 clerkships, indicating broad, distributed reinforcement across the clinical curriculum. In total, 1 (4%) MEPO had zero matching SLOs identified in any clerkship, a coverage gap that had not been previously identified through the institution's existing reporting processes. A total of 3 (12%) additional MEPOs were matched to SLOs in only a single clerkship, indicating narrow rather than absent coverage that may still warrant curricular attention. This analysis was limited to SLOs; assessment items had not yet been mapped to MEPOs in the clerkship database at the time of this report, consistent with the assessment-tagging extension described in the Future Directions section.

XSL•FO
RenderX

Table 3. Clerkship Medical Education Program Objective (MEPO) coverage report across all 9 clinical clerkships, generated via AI-agent query against the Airtable clerkship curriculum database, July 2026[a].

MEPO code	Matching clerkship SLOs[b], n	Clerkships covering (of 9)	Status
K1	97	9	Covered
K2	1	1	Single-clerkship only[c]
K3	97	9	Covered
K4	1	1	Single-clerkship only
K5	13	9	Covered
S1	16	8	Covered
S2	26	8	Covered
S3	14	8	Covered
S4	1	1	Single-clerkship only
S5	52	9	Covered
S6	97	9	Covered
S7	32	9	Covered
S8	13	9	Covered
S9	103	9	Covered
S10	22	7	Covered
S11	6	5	Covered
S12	11	9	Covered
S13	2	2	Covered
S14	6	5	Covered
AB1	23	9	Covered
AB2	10	8	Covered
AB3	0	0	Gap—no coverage[d]
AB5	11	9	Covered
AB6	13	9	Covered
AB7	22	9	Covered

[a]Full matched SLO codes by clerkship are available from the corresponding author.

[b]SLO: session learning objective.

[c]"Single-clerkship only" indicates the MEPO was matched to SLOs in exactly 1 of 9 clerkships.

[d]"Gap—no coverage" indicates no matching SLO was identified in any clerkship.

AI-Tagging Validation

Three independent course directors completed the blinded coding task; the number of items an individual rater excluded as not confidently codable from the SLO text alone ranged from 0 to 3 of 30. Agreement across the 3 raters and the database, summarized with Fleiss κ, was substantial at the exact-code level (κ=0.67, 95% CI 0.55-0.78) and almost perfect at the domain level (κ=0.94, 95% CI 0.88-1.00); the domain-level CI did not cross below the almost-perfect threshold at either bound. Pairwise comparison showed that this aggregate figure reflected 2 raters converging closely with each other and with the database (exact-code: κ=0.82-0.89; domain-level: κ=0.96-1.00) and a third rater whose exact-code agreement with the other 2 was lower (κ=0.47-0.55) despite comparable domain-level agreement (κ≥0.88), consistent with that rater's own report of difficulty completing portions of the task. Review of the residual exact-code disagreements showed most reflected genuine differences in expert framing at the subcode level for content plausibly spanning 2 adjacent categories (eg, whether a laboratory-test-interpretation objective is coded to the organ-system domain or to the biostatistics and epidemiology domain covering test characteristics) rather than database error; one item, on which the database initially diverged from all 3 raters, was corrected to match the converging rater consensus. This blinded, independent comparison is distinct from, and addresses the limitation of, the routine phase 2 verification described in the Methods section, which is performed by curriculum staff rather than independent raters and was not separately quantified.

XSL•FO
RenderX

Discussion

Principal Findings

These findings offer preliminary evidence that AI-integrated relational architecture can convert curriculum governance from a periodic, retrospective activity into a continuous, self-service one. The FOMI-to-ID vertical alignment analysis and the 16-course USMLE Step 1 report demonstrate that questions requiring days to weeks of manual staff effort can be answered within a single working day, and that resulting reports are specific enough to inform course-level action, including concrete recommendations such as session consolidation and content-ownership assignment, rather than generic classifications alone. The blinded validation study also offers an independent answer to a question the framework's design left open: how closely AI-proposed tags track subject-matter expert judgment once verified. Almost-perfect domain-level agreement (Fleiss κ=0.94) across 3 independent course directors and the database suggests the phase 2 workflow is not merely producing plausible-looking output that staff rubber-stamp, but output that substantively tracks how subject-matter experts would classify the same material; substantial, rather than almost-perfect, exact-code agreement (κ=0.67) is consistent with genuine ambiguity in content that spans 2 adjacent domains rather than with systematic AI error, a pattern reinforced by the fact that the 2 most consistent raters agreed with each other about as often as either agreed with the database. The narrower panel size (3 raters) means these κ estimates, while their CIs did not cross into lower agreement categories, should be interpreted as an initial rather than definitive estimate of population-level agreement. Regardless of this level of agreement, we want to be explicit that these findings are not a basis for reducing human review of AI-proposed codes. The validation study measures how closely a verified database matches independent expert judgment; it does not establish that any individual AI-proposed code can be trusted without the course-director confirmation step described in the Methods section, since even the raters who most closely agreed with the database and each other still diverged on a meaningful minority of items, and the courses represented in the discrepant items were not predictable in advance. We view the verification step as a permanent feature of this framework rather than a transitional safeguard to be relaxed as confidence in the underlying model grows.

The curriculum-wide pathology coverage deficit is consistent with the pattern described by Mahzari et al [2], in which foundational courses often underweight pathology relative to physiology and pharmacology. One plausible explanation is that faculty teaching in early preclerkship courses are often drawn from basic science rather than pathology backgrounds, and SLOs in these courses establish normal structure and function before disease processes are introduced elsewhere. However, the discipline-level performance comparison in the Results section complicates that a simple coverage explains performance account: physiology showed a below-national-mean performance gap nearly as large as pathology's despite coverage within its target range, so coverage alignment by itself does not fully explain relative discipline performance, and factors such as question difficulty, item exposure, or student study allocation likely also contribute. This descriptive association does not establish a causal or formally tested statistical relationship; a larger, prospective sample linking alignment scores to individual outcomes, controlling for other contributors such as prior preparation, is needed before stronger claims can be made. We return to this point in the Limitations section.

Comparison With Prior Work

This study builds directly on our previously published study describing the underlying USMLE coding library [8]. The present study reports the operational deployment of that library within a complete relational database and reporting framework, including faculty-facing search, cross-course redundancy detection, and accreditation-linked reporting not described in the earlier work. It also extends the theoretical foundation established by Harden [1] and the integration framework proposed by Youm et al [4] by operationalizing curriculum visibility as a queryable, faculty-facing technology rather than a static document. Consistent with the growing body of literature on large language models in medical education published in this journal, including a review of opportunities and challenges for large language models in medical education by Abd-Alrazaq et al [7], this study demonstrates a specific institutional governance application of generative AI that extends beyond the student-facing tutoring and content-generation applications more commonly reported to date. Unlike prior curriculum mapping efforts that relied on one-time computational analyses [6], the architecture described here is persistent operational infrastructure, generating new reports on demand as curriculum data evolve.

Positioned against a formal program-evaluation framework such as Kirkpatrick's 4-level model or the context, input, process, product model, the evidence reported here sits primarily at the process level: it demonstrates that the framework functions as designed, produces specific and actionable outputs, and yields tagging that agrees closely with independent expert judgment, but it does not yet include reaction-level data from faculty users, learning- or behavior-level data on whether curriculum decisions changed as a result of the reports, or product-level data linking the framework to downstream outcomes such as licensing examination performance. Kirkpatrick level 1 (reaction, via faculty usability data) and level 4 (results, via the planned outcome correlation described below) are both identified as priorities for subsequent work rather than claimed here.

Limitations

This study has several limitations. First, routine phase 2 verification of AI-proposed mappings by the Curriculum Mapping Unit was not accompanied by a formally logged acceptance/correction rate at the time of initial deployment; a retrospective, blinded validation study was subsequently conducted to address this gap and is reported in the Results section and discussed further below, but it postdates the original tagging decisions rather than having been built into the original workflow. That validation study also used a modest rater panel (3 course directors) and a single 30-item sample; the resulting κ estimates, while their CIs remained within their respective agreement categories, would benefit from a larger panel and sample in future work. Second, the pathology coverage and

XSL•FO
RenderX

Step 1 performance association is descriptive, not a formal statistical test controlling for other contributors to performance and should be read as hypothesis-generating rather than confirmatory. Third, the generative AI models used are proprietary and nondeterministic, constraining exact reproducibility; human verification mitigates but does not eliminate this constraint. Fourth, the framework was implemented at a single institution during a single AY, and findings on adoption, workload effects, and report quality may not generalize to institutions with different curricular structures, governance processes, or database platforms. Fifth, formal faculty usability and adoption data, such as search-query volume, satisfaction ratings, or structured usability testing, were not systematically collected during this reporting period; faculty engagement is described narratively in the Results section based on curriculum-committee and course-director feedback, and a structured usability survey is planned as near-term work. Finally, the accuracy of the searching and reporting depends on the conciseness and granularity of the underlying session and clerkship learning objectives; faculty are encouraged to use a staff-assisted AI prompt for writing granular objectives.

Future Directions

A near-term extension of the phase 2 tagging workflow will map all assessment items that are deidentified codes systematically to their corresponding SLOs, completing the bidirectional chain from examination item to SLO, CLO, MEPO, and USMLE content code across all preclerkship courses. Once in place, the AI agent will be able to generate blueprinting reports showing which program objectives any assessment evaluates, at what Bloom's Taxonomy level, and how that distribution aligns with USMLE discipline targets. This assessment-tagging extension had not yet been implemented at the time of this report and is planned as the focus of a subsequent publication.

Two further extensions are planned once assessment-item tagging is complete. First, codifying assessment items and linking them explicitly to their corresponding SLOs, described earlier, would allow both students and curriculum administrators, not faculty alone, to view which objectives are assessed, at what cognitive level, and how thoroughly, extending the framework's transparency function beyond its current faculty-only scope. Second, a student-facing version of the database is planned that would allow learners to search the curriculum by clinical concept and see, for each session, the corresponding USMLE content domain, providing the same transparent, self-directed curriculum visibility described for faculty to students preparing for licensing examinations.

A recurring challenge described by clerkship leadership at a peer institution illustrates the framework's potential for rapid-cycle improvement. In their current workflow, gap and redundancy reports are generated using an existing reporting system, but the same objective coverage issues recur across successive academic cycles because identified changes are not implemented before the clerkship resumes: with reporting turnaround measured in weeks, the narrow window, often as short as 2 weeks, between rotation blocks in which curricular revisions can practically be made closes before the relevant report becomes available. Because the present framework's AI-agent reporting layer returns gap and redundancy analyses within minutes rather than weeks, a report generated at the close of one rotation block can be reviewed and acted on well within the following block's revision window. The same leadership described a related diagnostic gap: their system can confirm that a fixed number of institutional program objectives are formally defined but can successfully trace curricular coverage for only a subset of them, with no mechanism to identify which objective is missing or why. Because the present framework tags every SLO explicitly to specific MEPOs at the point of data entry, an AI-agent query can trace the chain of linked SLOs to determine where in curriculum design a gap originates. This diagnostic capability was demonstrated directly within our own institution's clerkship phase (see Clerkship MEPO Coverage Analysis section; Table 3), where a single query traced all 25 MEPOs against clerkship SLOs across 9 rotations and identified 1 MEPO with no matching SLO in any clerkship, a gap not previously visible through existing reporting processes. Extending this tagging discipline to assessment items would enable the equivalent diagnostic capability for assessment coverage specifically.

The most institutionally significant planned application of this infrastructure is expected to be its support for LCME accreditation preparation, particularly for Standard 6, which requires schools to document that MEPOs are reflected throughout the curriculum and demonstrably map to curricular content and, eventually, to assessments. Airtable's native Claude AI integration offers a no-configuration pathway to the reporting capabilities described here, and future work will evaluate whether this native integration reproduces the analytical quality of structured-prompt API methods. This framework can be built at any medical school regardless of size or infrastructure, requiring no proprietary curriculum management platform and no institution-level AI servers; ongoing costs, limited to a cloud database subscription and AI-agent query fees, typically total under several hundred dollars per AY at current pricing. This cost structure is a source of risk as well as accessibility: because the framework depends on 2 external commercial services rather than institution-owned infrastructure, a substantial vendor price increase, a change in API terms, or discontinuation of either service could require migration effort or increased recurring cost. The framework's database-agnostic and model-agnostic design, described in the Methods section, is intended to mitigate this risk by allowing either component to be replaced without redesigning the underlying data structure, but this substitutability has not been tested in practice, and institutions adopting this approach should budget for the possibility of a migration.

The present framework was built around the USMLE Content Outline and LCME accreditation standards, which apply specifically to US MD-granting programs. Adaptation to US osteopathic medical education would require substituting a Comprehensive Osteopathic Medical Licensing Examination (COMLEX)–aligned content outline for the USMLE library described here; because COMLEX and USMLE differ in content organization and osteopathic-specific domains, this would require rebuilding the coding library rather than simply relabeling the existing one, though the relational database

XSL•FO
RenderX

architecture and AI-tagging workflow themselves would not need to change. International medical schools present a larger adaptation challenge, since many national or regional licensing examinations do not publish a content outline with the same structure, granularity, or public availability as the USMLE Content Outline, and some jurisdictions have no equivalent standardized examination at all; institutions in those settings would likely need to construct a comparable local coding library from their own accreditation or licensing body's standards before this framework's tagging and reporting methodology could be applied. Adaptation for nursing, pharmacy, dental, and postgraduate training is comparatively more straightforward wherever a curriculum aligns to a standardized external examination with a published content outline. We have not tested the framework in any of these adjacent contexts and present this discussion as an assessment of adaptation requirements rather than a demonstrated capability.

Conclusions

An AI-integrated relational curriculum database with a faculty-facing search interface and an AI-agent reporting layer can make curriculum content, redundancy, and USMLE alignment visible and actionable on a timescale of minutes rather than weeks, producing specific recommendations rather than aggregate metrics alone. This implementation study reports the first phase of a larger program of work; extension to assessment-item tagging and prospective validation against licensing examination outcomes are planned as subsequent contributions.

Acknowledgments

The authors thank Fabiola Saldivar, MD, MSHS; Anthony Hackett, DO, FACEP, FAAEM, FRSM; and Anitha Shenoy, PhD for independently coding session learning objectives against the USMLE (United States Medical Licensing Examination) Content Outline as part of the blinded validation study, contributing the rater data underlying the interrater reliability analysis reported in the Results section. The authors also thank Bridget Lighthall, MEd; Dustin Dubois, PhD; and the Office of Academic Affairs team for their support and encouragement. Anthropic Claude Sonnet 5, accessed via API, was used as described in the Methods section to propose USMLE content code, Association of American Medical Colleges (AAMC) keyword, and Bloom's Taxonomy tags for session learning objectives; all AI-generated tags were reviewed and verified by the Curriculum Mapping Unit before database entry. An AI agent connected to the study database via API was used to generate the alignment, redundancy, gap, and accreditation reports described in the Results section. Elements of this work were presented at the Southern Group on Educational Affairs 2026 Annual Meeting and submitted to the AAMC 2026 Staff Track Call for Proposals.

Funding

The authors declared no financial support was received for this work.

Data Availability

The USMLE (United States Medical Licensing Examination) Content Outline coding library and the Airtable database template generated during this study are available from the corresponding author on reasonable request.

Generative AI Use

The authors declare that generative AI was not used in the preparation of this manuscript.

Authors' Contributions

RMB wrote the manuscript and developed the innovation. JL assisted RB with data collection and processing. DD provided senior supervision and oversight.

Conflicts of Interest

None declared.

Multimedia Appendix 1

Representative prompt templates and classification decision rules used for phase 2 AI-agent mapping of session learning objectives to USMLE (United States Medical Licensing Examination) Content Outline codes, including the exact prompt structure, the rules governing code selection when an objective could plausibly map to more than 1 domain, and example inputs and outputs.
[PDF File (Adobe PDF File), 166 KB-Multimedia Appendix 1]

Multimedia Appendix 2

Screenshot of the faculty-facing read-only Airtable clerkship curriculum database (AY25_26_Clerkship_Map_USMLE_NBME), shown for the Emergency Medicine (EM) rotation (academic year 2025-2026). Rotation tabs and a dedicated Medical Education Program Objective (MEPO) Reference tab provide direct access to all clerkship disciplines. Each record links a clerkship core

XSL•FO
RenderX

objective to its rotation-specific session learning objective, mapped MEPO, United States Medical Licensing Examination Step 2 content topic, and corresponding National Board of Medical Examiners (NBME) content category, extending the preclerkship database architecture to include direct NBME mapping.
[PNG File , 162 KB-Multimedia Appendix 2]

References

1. Harden R. AMEE Guide No. 21: curriculum mapping: a tool for transparent and authentic teaching and learning. Med Teach. 2001;23(2):123-137. [doi: 10.1080/01421590120036547] [Medline: 11371288]
2. Mahzari M, AlNahedh T, Ahmed AA, Al Rumyyan A, Shaban S, Magzoub ME. Practical guide to undergraduate medical curriculum alignment and mapping. Adv Med Educ Pract. 2023;14:1001-1012. [FREE Full text] [doi: 10.2147/AMEP.S424815] [Medline: 37745032]
3. Barde PB, Parmar N, Chitturi V, Sharma G, Kathrotia R, Joshi K, et al. Curricular redundancy in medical undergraduate course: critical insights from various stakeholders in medical institutes in India. Avicenna J Med. 2024;14(3):167-174. [FREE Full text] [doi: 10.1055/s-0044-1791843] [Medline: 39584165]
4. Youm J, Christner J, Hittle K, Ko P, Stone C, Blood AD, et al. The 6 degrees of curriculum integration in medical education in the United States. J Educ Eval Health Prof. 2024;21:15. [FREE Full text] [doi: 10.3352/jeehp.2024.21.15] [Medline: 38872249]
5. Al-Eyd G, Achike F, Agarwal M, Atamna H, Atapattu DN, Castro L, et al. Curriculum mapping as a tool to facilitate curriculum development: a new school of medicine experience. BMC Med Educ. 2018;18(1):185. [FREE Full text] [doi: 10.1186/s12909-018-1289-9] [Medline: 30081897]
6. Takizawa PA. Using a topic model to map and analyze a large curriculum. PLoS One. 2023;18(4):e0284513. [FREE Full text] [doi: 10.1371/journal.pone.0284513] [Medline: 37079546]
7. Abd-Alrazaq A, AlSaad R, Alhuwail D, Ahmed A, Healy PM, Latifi S, et al. Large language models in medical education: opportunities, challenges, and future directions. JMIR Med Educ. 2023;9:e48291. [FREE Full text] [doi: 10.2196/48291] [Medline: 37261894]
8. Badeau RM, Mallini T, Laury J, VanBuren V, Dickey D. From parallel to aligned: a United States Medical Licensing Examination (USMLE) mapping and assessment-tagging framework for medical education. BMC Med Educ. 2026;26(1):1606. [FREE Full text] [doi: 10.1186/s12909-026-09912-2] [Medline: 42449355]
9. Biggs J, Tang C. Teaching for Quality Learning at University. 4th Edition. London, United Kingdom. Open University Press; 2011.
10. Wijnen-Meijer M, van den Broek S, Koens F, Ten Cate O. Vertical integration in medical education: the broader perspective. BMC Med Educ. 2020;20(1):509. [FREE Full text] [doi: 10.1186/s12909-020-02433-6] [Medline: 33317495]
11. Title 45, Part 46, Section 46.102: Definitions for purposes of this policy. Code of Federal Regulations. URL: https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-A/part-46/subpart-A/section-46.102 [accessed 2026-10-29]
12. Functions and structure of a medical school. Liaison Committee on Medical Education. 2026. URL: https://lcme.org/standards-and-rules/ [accessed 2026-10-22]

Abbreviations

AAMC: Association of American Medical Colleges
AY: academic year
COMLEX: Comprehensive Osteopathic Medical Licensing Examination
CLO: course learning objective
FERPA: Family Educational Rights and Privacy Act
FOMI: Foundations of Medicine I
FOMII: Foundations of Medicine II
ID: Introduction to Disease
LCME: Liaison Committee on Medical Education
MEPO: Medical Education Program Objective
POMI: Practice of Medicine I
POMII: Practice of Medicine II
SLO: session learning objective
USMLE: United States Medical Licensing Examination

XSL•FO
RenderX

Edited by D Chartash; submitted 22.Jul.2026; peer-reviewed by SSR Anugu, A M, Y Lim; comments to author 10.Sep.2026; revised version received 25.Sep.2026; accepted 14.Oct.2026; published 12.Nov.2026

Please cite as:
Badeau RM, Laury J, Dickey D
An AI-Integrated Relational Curriculum Mapping Framework for Alignment and Transparency in Medical Education: Implementation Study
JMIR Med Educ 2026;12:e107727
URL: https://mededu.jmir.org/2026/1/e107727
doi: 10.2196/107727
PMID:

©Robert M Badeau, Jennifer Laury, Danielle Dickey. Originally published in JMIR Medical Education (https://mededu.jmir.org), 12.Nov.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR Medical Education, is properly cited. The complete bibliographic information, a link to the original publication on https://mededu.jmir.org/, as well as this copyright and license information must be included.

XSL•FO
RenderX