

SyllaMatch: An LLM Powered Credit Transfer System using Multi-Agent Workflow

Kazim Timucin UTKAN
Industrial Engineering
Istanbul Technical University
Istanbul, Türkiye
timucinutkan@gmail.com

Elifnaz OLGAC
Industrial Engineering
Istanbul Technical University
Istanbul, Türkiye
enazolgac@gmail.com

Oguzhan KAHRAMAN
Data Engineering and Business Analytics
Istanbul Technical University
Istanbul, Türkiye
oguzhan2639@gmail.com

Abstract—Matching syllabi for inter-institutional transfer students is a significant burden on the academic personnel and can strongly shape students’ academic path. This process of aligning past coursework to the receiving institution’s course catalog is a tedious, mostly manual and time-consuming process. Academic staff responsible for this workflow processes high-volume data within tight time constraints and the resulting decisions often lack transparency. This limited transparency damages students’ trust and commitment to the new institution. This project proposes SyllaMatch, a Large Language Model (LLM) based automated decision support system for the syllabus matching process of transfer students. The system provides an end-to-end pipeline which encompasses data extraction from student transcripts and course catalogs, utilizing LLM to extract a structured list of courses and Retrieval-Augmented Generation (RAG) framework to perform a semantic matching of each course against the receiving institution’s catalog. Finally, the system produces a comprehensive exemption report that presents course match suggestions with clear evidence and justifications. By providing consistency, traceability and efficiency, SyllaMatch supports a data-oriented approach to course matching process, leaving academic staff responsible only for reviewing and final approval.

Index Terms—syllabus matching, course equivalency system, large language models, retrieval-augmented generation, decision support system.

I. INTRODUCTION

Undergraduate education is regarded as a fundamental part of contemporary life. For this reason, a large population of youth pursues undergraduate degrees as a standard step within their lifespan. Yet, students’ undergraduate paths may sometimes be affected by various factors. These factors may include relocation, financial constraints, cultural and social adjustment challenges, leading to a need for institutional transfer of the student.

In this case, preserving students’ previous academic workload, completed in the sending institution, requires conducting a systematic course matching process. In course matching process, the main expectation of the students is that their earned credits to be recognized to the greatest extent possible. Whereas the receiving institution seeks to ensure that the incoming student’s academic background aligns with the learning outcomes and requirements of their programs’ current curriculum. These differences in primary objectives may lead to disagreements or perceived unfairness among students

regarding the decision-making process. The disagreement is further intensified by the limited transparency in the decision process, as although the academic personnel conducts credit matching in accordance with the institutional guidelines, the final decision is made by a human evaluator and therefore may be perceived as subjective.

The course exemption process also varies across different institution types. In public universities, the main challenge is that a high volume of transfer applications must be evaluated within a limited time frame. This workflow includes labor-intensive evaluation of heterogeneous documents such as, syllabi, transcripts, course catalogs and institutional guidelines for the equivalency rules. This proves a heavy administrative load, additionally, the exemption process is usually accomplished in a limited time set by the institution. In contrast, private universities have different priorities regarding students transfers. While the time frame set by the institution may be more flexible, the response time to students becomes a crucial parameter. As is often the case, the candidate students apply in multiple schools simultaneously, and the number of credits they can transfer becomes a key information for their enrollment decision. Therefore, a well-communicated and an accelerated decision-making process for course exemption process provides a competitive advantage for private institutions.

Considering both these cases, although the institution type changes the accreditation priorities and constraints, the common need for the course transfer process remains to be reducing workload, improving transparency, supporting consistent and evidence-based decision making.

SyllaMatch proposed in this study, offers an LLM powered decision-making system that encompasses extracting structured information from documents, matching course similarity based on semantic level using Retrieval-Augmented Generation (RAG) framework and providing evidence behind course matches on a final exemption report. Integration of SyllaMatch into the course equivalency process significantly reduces the time required for course matching and leaves academic personnel responsible only for the final review and approval of the exemption report.

II. RELATED WORK

There is limited work on course transfer in literature. Chandrasekaran and Mago developed a clustering-inspired system that compares structured text inputs of learning outcomes to evaluate the similarity of courses [1]. Lisboa et al. developed a natural language processing (NLP) based model to analyze course contents and provide matching courses specifically tailored for Portuguese education systems [2]. Pardos et al. focus on the number of credits transferred through programs by making predictions using machine learning models [3].

Additionally, there are several tools for automating credit transfer systems in the market. Automated Transfer Equivalency system from University of Texas is a database containing syllabus matches from prior years [4]. CollegeSource Transfer Evaluation system provides an interactive catalog of the institutions courses and the historical matches [5]. DegreeSight is a commercial artificial intelligence (AI) based similarity matching product, available only for U.S. colleges and universities [6].

While these approaches present important development towards automating syllabus matching process, most of the existing tools rely heavily on text similarity and existing databases of historical databases. These dependencies limit the applicability of such tools for different institutions. Additionally, the proposed models focus either on the course outcomes or context, lacking a holistic view and a semantic grounding while deciding on equivalency evaluations. Hence, there is a requirement for a flexible course transfer application that offers flexible, interpretable, and transparent system that enables an end-to-end approach to course matching process.

III. SYSTEM ARCHITECTURE

The architectural design of SyllaMatch marks a departure from traditional linear conversational agents. Instead of a simple input-output mechanism, the system is engineered as a **Directed Cyclic Graph (DCG)**. This topological choice models the complex, non-deterministic nature of academic decision-making, allowing the system to execute iterative loops, recover from errors, and seamlessly integrate human oversight when necessary.

A. Architectural Paradigm: Agentic Workflow

The core operational logic is governed by a state machine comprising 14 distinct functional nodes. Unlike linear chains (DAGs) commonly used in basic LLM applications, this cyclic architecture enables the "ReAct" (Reasoning + Acting) pattern. For instance, if the system detects missing syllabus information, it does not simply fail or hallucinate; it transitions into a search loop, autonomously querying external sources until the data deficit is resolved or a retry limit is reached.

The workflow is divided into four high-level phases:

- 1) **Parallel Parsing & Ingestion:** The system processes PDF transcripts and course content files simultaneously. It utilizes `markitdown` for deep text extraction, with a fallback to Vision models for complex layouts, ensuring high-fidelity data ingestion from heterogeneous formats.

- 2) **Structured Extraction:** Raw text is converted into strict JSON objects using Pydantic schemas. This ensures that critical entities such as Course Code, ECTS credits, and Grades are extracted with type safety, mitigating the risk of LLM hallucinations.
- 3) **Autonomous Resolution (HITL):** A "Check Content" node audits the extracted data. If a course description is missing, the workflow pauses, triggering a **Human-in-the-Loop (HITL)** state. The academic staff can then direct the **AI Search Agent** to scrape specific university domains (e.g., .edu, .edu.tr) via the Tavily API to retrieve the missing syllabus.
- 4) **Semantic Matching Loop:** The system iterates through each valid student course, retrieving the top 5 candidates from the target curriculum using vector similarity and passing them to the "Judge" node for final adjudication.

B. Technology Stack and Selection Rationale

The technology stack was rigorously selected to prioritize reliability, observability, and the handling of non-deterministic AI behaviors. The rationale for key components is as follows:

- 1) **Orchestration: LangGraph:** LangGraph was selected over standard orchestration libraries due to its native support for cyclic graphs and state persistence. This capability is essential for managing long-running agentic workflows where the system state (e.g., the list of processed courses) must be preserved across user interruptions and search iterations.

- 2) **Inference Engine: GPT-5.2:** The system utilizes **GPT-5.2** as the primary reasoning engine. Academic evaluation requires a nuanced understanding of learning outcomes that smaller or older models often fail to capture. GPT-5.2 is specifically employed in the "Judge" phase to analyze the semantic overlap between source and target courses, providing reasoning capabilities superior to standard similarity metrics.

- 3) **Vector Database: Supabase & pgvector:** We utilize Supabase (PostgreSQL) extended with `pgvector` to provide a unified platform for both relational data (student requests) and high-dimensional vector storage. This simplifies the infrastructure by eliminating the need for a separate vector database. The system uses Hierarchical Navigable Small World (HNSW) indexes to perform fast approximate nearest neighbor searches on 1536-dimensional embedding vectors.

- 4) **Search Infrastructure: Tavily API:** Standard search APIs often return cluttered HTML that distracts LLMs. The Tavily API was chosen for its ability to return parsed, clean, and execution-ready content specifically optimized for LLM context windows. Furthermore, its strict domain filtering capabilities ensure that the agent only retrieves data from trusted academic sources, reducing the risk of false information.

C. Embedding Strategy

To capture the true semantic meaning of a course beyond its title, SyllaMatch employs a **Weighted Average** embedding strategy. The system separately embeds the Course Title, Objectives, Course Content, and Learning Outcomes using the

text-embedding-3-large model. These individual vectors are then averaged—potentially with learnable weights—to create a single, rich composite vector. This ensures that a course titled "Intro to Coding" can be accurately matched with "Computer Programming I" based on the underlying concepts (e.g., loops, arrays) rather than mere keyword overlap.

IV. METHODOLOGY

The operational workflow of SyllaMatch is designed to simulate the cognitive process of a human academic evaluator, enhanced by the speed and consistency of generative AI. The methodology follows a rigorous four-stage pipeline: Ingestion, Resolution, Retrieval, and Adjudication.

A. Phase 1: Multi-Modal Ingestion and Extraction

The process begins with the ingestion of the student's official transcript and any accompanying course content files (PDFs). Unlike standard OCR tools that often lose structural context, SyllaMatch employs a conditional parsing strategy:

- 1) **Deep Text Parsing:** The system initially attempts to parse the document using `markdown`, aiming to reconstruct the text while preserving table structures and headers.
- 2) **Vision Fallback:** If the extracted text coherence score falls below a predefined threshold (indicating complex layouts or scanned images), the system triggers a fallback to the GPT-5.2 Vision model. This multi-modal approach ensures high-fidelity extraction even from legacy or poorly scanned documents.

Following raw text extraction, the `extract_courses` node utilizes Large Language Models (LLMs) constrained by Pydantic schemas. This enforces strict output formatting, ensuring that entities such as "Course Code," "ECTS Credit," and "Grade" are extracted as typed JSON objects rather than unstructured text.

B. Phase 2: Agentic Content Resolution (HITL)

A unique contribution of this study is the handling of data scarcity. Transcripts often list course titles (e.g., "CS101") without the detailed syllabus required for evaluation. SyllaMatch addresses this via a **Human-in-the-Loop (HITL)** architecture:

- **Detection:** The `check_content` node scans all extracted courses. If a course lacks a detailed description, it flags the item as "Content Missing."
- **Intervention:** The workflow suspends execution and presents the user with a decision matrix: (a) Manually upload content, (b) Generate synthetic summaries (not recommended for final decisions), or (c) Activate the **AI Search Agent**.
- **Autonomous Search:** If the Search Agent is activated, it initiates a "ReAct" (Reasoning + Acting) loop using the Tavily API. The agent formulates queries (e.g., "Technical University A CS101 syllabus learning outcomes"), filters results for academic domains (e.g., .edu), scrapes the syllabus, and verifies its relevance before injecting it back into the workflow.

C. Phase 3: Semantic Retrieval (RAG)

Once all student courses have valid descriptions, the system moves to the matching phase. We employ a Retrieval-Augmented Generation (RAG) framework utilizing the Supabase vector store. For each student course C_s , the system generates a query vector V_{C_s} using the weighted average strategy described in Section III. It then performs a cosine similarity search against the target curriculum index to retrieve the top $k = 5$ candidate courses ($C_{t1} \dots C_{t5}$).

D. Phase 4: The LLM Judge (Adjudication)

The final and most critical phase is the **LLM Judge**. Standard RAG applications often rely solely on similarity scores, which can be misleading (e.g., "Advanced Physics" and "Intro to Physics" may have high vector similarity but are not equivalent). SyllaMatch introduces a discrete `judge` node powered by GPT-5.2. This node acts as an academic arbitrator. It takes the student's course content and the top retrieved target course as input and performs a comparative analysis of their **Learning Outcomes**.

Decision Logic Example: In a case study involving "Introduction to Programming," the RAG system retrieved "CMPE101" as a high-similarity candidate. However, the Judge node detected that the student's course focused on Python and high-level abstractions, whereas the target CMPE101 required strict knowledge of C++ memory management and pointers. Consequently, despite the high vector similarity score, the Judge output a **REJECTED** decision with the explicit reasoning: "Target course requires pointer arithmetic and memory management competencies absent from the source syllabus".

V. EVALUATION

In the domain of academic equivalency, traditional machine learning performance metrics such as Accuracy or F1-Score often fail to capture the nuanced reality of institutional decision-making. Syllabus matching is inherently subjective, governed by specific university regulations and the discretionary judgment of academic commissions. Therefore, our evaluation framework prioritizes *Justification Quality* and *Decision Support Reliability* over binary classification success.

A. Expert-in-the-Loop Validation Protocol

To validate the model, we designed a Human-in-the-Loop verification protocol. The system processes pairs of syllabi and generates a comprehensive report detailing:

- **Thematic Alignment:** Overlap in weekly course topics and learning outcomes.
- **Structural Compatibility:** Comparison of ECTS credits, contact hours, and course levels.
- **Justification Narrative:** A natural language explanation justifying the suggested equivalency (or rejection).

Subject Matter Experts (SMEs) audit these outputs against the original source documents, rating each justification based on whether the LLM correctly identified critical thematic pillars and adhered to the institution's specific transfer credit policies.

B. Quantitative Metrics

To quantify this qualitative feedback, we utilize two primary metrics:

- 1) **Justification Agreement Rate (JAR):** Unlike simple accuracy, this measures the percentage of model-generated *explanations* deemed sufficient and accurate by academic evaluators.
- 2) **Inter-Rater Reliability (Cohen’s Kappa):** To ensure the system remains objective across different evaluators, we calculate the Kappa Score (κ). This measures the degree of consensus between multiple experts regarding the model’s justifications, adjusting for the possibility of agreement occurring by chance.

A high Kappa coefficient demonstrates that SyllaMatch provides a standardized baseline for equivalency that transcends individual faculty bias.

C. Continuous Improvement

Recognizing that receiving institutions have varying thresholds for credit transfer, the model is designed for contextual adaptation. The evaluation phase includes iterative feedback loops via structured surveys to refine the prompt engineering and retrieval mechanisms, ensuring the system evolves from a general matching tool into a specialized decision-support mechanism tailored to the host university.

VI. CONCLUSION

This study presented **SyllaMatch**, an agentic AI system designed to automate and standardize the university course transfer process. By integrating the reasoning capabilities of **GPT-5.2** with a semantic RAG framework, the system addresses the critical challenges of administrative workload, process opacity, and student dissatisfaction.

The proposed architecture moves beyond the limitations of keyword-based matching and static historical databases. Through its *Directed Cyclic Graph* workflow, SyllaMatch successfully handles data scarcity by autonomously retrieving missing syllabus information and employs an “LLM Judge” to provide evidence-based adjudication.

Our evaluation methodology highlights that successful automation in this domain requires more than just high retrieval accuracy; it demands transparent, interpretable justifications that align with institutional policies. SyllaMatch achieves this by providing consistent, traceable, and efficient decision support, ultimately leaving academic personnel responsible only for the final review and approval, thereby transforming a tedious administrative burden into a streamlined, data-driven process.

Future work will focus on expanding the curriculum database to include a broader range of international institutions and further refining the autonomous agent’s ability to navigate complex university authentication systems.

REFERENCES

- [1] D. Chandrasekaran and V. Mago, “Automating Transfer Credit Assessment-A Natural Language Processing-Based Approach,” *Computers, Materials and Continua*, vol. 73, no. 2, pp. 2257–2274, 2022.
- [2] M. Lisboa, D. Da Silva Passos, D. Nadler Prata, L. Eduardo Bovolato, and D. Paixão Pinheiro, “A Support Tool to Improve Course Credit Transfer in an Education Institution,” *DESAFIOS*, vol. 9, no. 2, pp. 159–178, 2022.
- [3] Z. A. Pardos, H. Chau, and H. Zhao, “Data-assistive course-to-course articulation using machine translation,” in *Proceedings of the 6th 2019 ACM Conference on Learning at Scale*, 2019.
- [4] “Transfer Credit Resources - University of Texas Admissions.” [Online]. Available: <https://admissions.utexas.edu/apply/transfer-students/transfer-credit-resources/>
- [5] “Transfer tools – CollegeSource.” [Online]. Available: <https://collegesource.com/transfer-tools/>
- [6] “Transfer Credit Evaluation Software — DegreeSight.” [Online]. Available: <https://www.degreesight.com/>