Journal of Information Systems Education

Volume 37
Issue 4
Fall 2026

Supporting Universities in Keeping Up With Technological Change: A GenAI Tool for Curriculum Analysis

Luis SaCouto, Tan Gürpinar, Shizhen (Jasper) Jia, and Sana Quadri

Recommended Citation: SaCouto, L., Gürpinar, T., Jia, S., & Quadri, S. (2026). Supporting Universities in Keeping Up With Technological Change: A GenAI Tool for Curriculum Analysis. *Journal of Information Systems Education*, 37(4), 625-638. https://doi.org/10.62273/RATM2670

Article Link: https://jise.org/Volume37/n4/JISE2026v37n4pp625-638.html

Received:	January 14, 2026
First Decision:	January 18, 2026
Accepted:	February 8, 2026
Published:	December 15, 2026

Find archived papers, submission instructions, terms of use, and much more at the JISE website:
https://jise.org

ISSN: 2574-3872 (Online) 1055-3096 (Print)

Journal of Information Systems Education, 37(4), 625-638, Fall 2026
https://doi.org/10.62273/RATM2670

Supporting Universities in Keeping Up With Technological Change: A GenAI Tool for Curriculum Analysis

Luis SaCouto
Tan Gürpinar
Shizhen (Jasper) Jia
Sana Quadri
School of Business
Quinnipiac University
Hamden, CT 06518, USA
Luis.SaCouto@qu.edu, Tan.Gurpinar@qu.edu, Shizhen.Jia@quinnipiac.edu, Sana.Quadri@qu.edu

ABSTRACT

In an era of rapid technological change, universities must continually adapt their curricula to prepare students for evolving careers. Yet systematic analysis and benchmarking of university course offerings across institutions remain labor-intensive and slow, limiting timely, data-informed curriculum decisions. This paper presents a GenAI-enabled tool for curriculum analysis that automates the extraction, retrieval, and categorization of course information from university course catalogs, enabling scalable and reusable analysis across topics and institutions. The tool integrates search with generative AI to identify relevant courses and classify them along user-defined dimensions (e.g., topic, use case, discipline), substantially reducing the manual effort required for cross-institutional curriculum comparison. We demonstrate the tool's capabilities through a case study of extended reality (XR) offerings in AACSB-accredited business schools, showing that it can replicate and extend prior manual curriculum studies with high accuracy while operating at significantly greater scale and speed. Designed for IS and computing education stakeholders, the tool supports curriculum design, academic benchmarking, and alignment with emerging industry skill demands. By transforming distributed and unstructured course catalog descriptions into structured, comparable insights, this work provides a practical, generalizable solution to help universities keep pace with technological change through evidence-based curriculum analysis.

Keywords: Generative AI, IS curriculum, Workforce development, Curriculum design & development, Industry-relevant skills, Program assessment & design

1. INTRODUCTION

In today's rapidly evolving digital landscape, the acceleration of technological innovation is reshaping industries and redefining the skills required for the modern workforce. Emerging technologies such as artificial intelligence (AI), extended reality (XR), distributed ledger technologies (DLT), and quantum computing (QC) are advancing rapidly and converging (Düdder et al., 2021). This leads to shorter innovation cycles and increased pressure on organizations to adapt (Gürpinar, 2025).

This dynamic environment necessitates that educational institutions, particularly universities, continuously update and align their curricula with industry demands. This puts pressure on faculty and

program coordinators to regularly revisit and realign course content with emerging competencies (Jovanović et al., 2025). Leading universities have begun integrating emerging technologies across disciplines rather than treating them as isolated topics (Awwad, 2024). The World Economic Forum projects that by 2025, half of all employees will require reskilling, highlighting the urgency for educational institutions to stay on track with their curricula (Hoteit et al., 2023). However, traditional methods of curriculum analysis and benchmarking are often manual, time-consuming, and inconsistent, making it challenging to keep pace with the rapid changes of our time (Jovanović et al., 2025).

Large language models (LLMs) have emerged as powerful tools capable of processing and understanding vast amounts of unstructured textual data. Their ability to perform semantic analysis, extract meaningful patterns, and classify information makes them well-suited for tasks that have traditionally required manual interpretation. In the context of fast-changing courses in higher education, these capabilities open the door to a new way to gain up-to-date insights. In the IS and computing education community, benchmarking curriculum content across institutions is critical for ensuring programs remain current (Ceccucci et al., 2020; Gürpinar et al., 2025a). This study addresses the scalability challenge of such benchmarking through automation. Specifically, this paper addresses the question of *whether generative AI can be reliably used to automate the large-scale extraction and analysis of academic course offerings* – a task traditionally done manually. We propose an LLM-based system to systematically analyze university course catalogs at scale through three core components: the extraction of course data from university websites, the filtering of relevant courses based on user-defined interests, and the categorization of courses along tailored curricular dimensions.

To explore this approach, the paper is structured as follows: Section 2 reviews related work, including manual curriculum benchmarking, existing course aggregation tools, and recent developments in GenAI-driven classification methods. Section 3 outlines our methodology, which includes extracting human-labeled ground-truth data and standard accuracy, precision, and recall metrics. Section 4 presents the architecture and workflow of our proposed system, detailing its components for extraction, filtering, and categorization. Section 5 showcases a case study focused on extended reality (XR) course offerings in U.S. universities with AACSB-accredited business schools, demonstrating the system's ability to replicate and extend a manually conducted analysis. Finally, Section 6 discusses key findings, limitations, and opportunities for future development, such as expanding the system to multilingual catalogs or integrating an interactive dashboard.

2. RELATED WORK

Curriculum benchmarking has traditionally been done through painstaking manual review. The analysis of academic curricula via manual methodologies presents significant limitations in scalability, consistency, and adaptability. For instance, the studies by Ceccucci et al. (2020) and Gürpinar et al. (2025a) involved a manual review of over 500 university programs to assess topics covered in business analytics curricula, or determine the integration of blockchain and web3 topics. Similarly, the research by Grimus (2020) and Busulwa et al. (2025) conducted a detailed examination of different curricula that evaluated the incorporation of emerging technologies. While these studies provide valuable insights, their manual and time-consuming nature makes them less responsive to the rapidly evolving educational landscape. Initial attempts at automated curriculum analysis have shown promise. Orellana et al. (2018), for instance, used text mining term frequency–inverse document frequency (TF-IDF) to assess course syllabi against a computing curriculum framework. Their findings highlighted the potential of text-based methods for comparing curricula across universities.

Recent research has explored the promising application of LLMs in various educational contexts. For example, Chan and Li (2025) developed a multimodal chatbot leveraging ChatGPT and Google Bard to enhance personalized learning experiences; Dickey and Bejarano (2024) introduced the GAIDE framework utilizing generative AI to assist in course content development. In the realm of analyzing educational content and providing automated feedback, Naseer et al. (2024) examined the role of generative AI in enhancing the scalability of in-class assessment processes. Beyond classroom use, AI has begun to assist

in curriculum design and evaluation tasks. For example, Zaki et al. (2023) automated the mapping of course learning outcomes to program outcomes using NLP, achieving ~85% accuracy compared to expert mapping. Similarly, Zamecnik et al. (2024) showed that LLMs can help verify whether curricula cover expected employability skills, with LLM-based skill annotations largely agreeing with human experts. These studies underline AI's growing role in curriculum analytics and quality assurance.

Despite these advancements, there remains a gap in the application of LLMs for the automated extraction, filtering, and categorization of university course offerings at scale. Existing studies have primarily focused on class content generation, personalized learning, or assignment assessment, i.e., limited exploration into academic catalogs. This paper aims to bridge this gap by proposing a comprehensive LLM-based system designed to automate the analysis of university course catalogs. No prior work has leveraged LLMs to automatically extract and classify entire university course catalogs for curriculum intelligence. This project addresses that specific gap.

3. RESEARCH METHODOLOGY

Our research methodology consists of three phases: (1) developing the GenAI-driven system, (2) preparing a ground-truth dataset for XR courses, and (3) evaluating the system's performance against this ground truth.

3.1 System Development

The core contribution of this work is the design and implementation of a modular pipeline that leverages large language models (LLMs) for curriculum intelligence. The system was developed iteratively, informed by practical requirements derived from a previous manually coded study (Gürpinar et al. 2025b). Its goal is to replicate and scale those processes. Specifically, the system follows a funnel-like structure (see Figure 1): it begins by collecting course data from a list of universities, then filters those courses for relevance to the topic under analysis, and finally categorizes them along key dimensions.

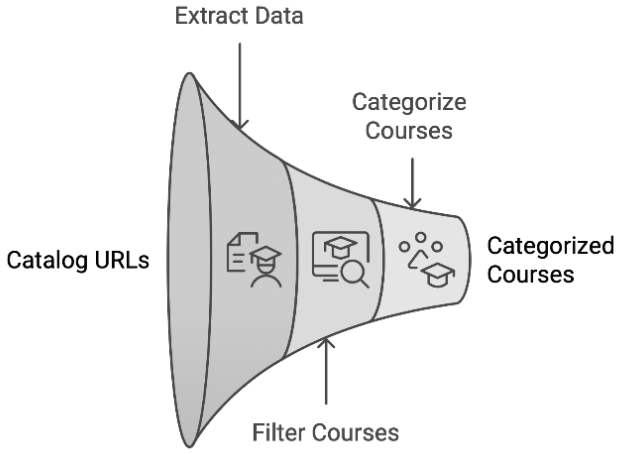


Figure 1. Research Process

During system development, we prioritized:

- *Automation:* replacing manual search and classification with LLM- and rule-based automation.
- *Flexibility:* enabling users to define arbitrary curricular interests and classification schemes.
- *Modularity:* allowing components (extraction, filtering, categorization) to operate independently or as part of an end-to-end pipeline.

Technically, the prototype was implemented in Python using scraping utilities, semantic search libraries, and OpenAI's GPT models accessed via API. A detailed description of the system follows in the next section.

3.2 Dataset Preparation

We focused on XR as an exemplar emerging technology domain, given its increasing relevance in IS and computing programs and prior interest in mapping XR education. From there, we drew on a dataset of publicly available course catalogs from over 500 U.S.-based institutions, with a focus on AACSB-accredited universities to maintain continuity with related studies.

From these catalogs, we manually constructed:

- A sample of courses from over 50 universities, where for each we gathered title, description, category, university, and level.
- A filtered set with all XR-relevant courses from over 50 universities, labeled by domain experts based on shared interest criteria.
- A multi-label classification dataset of over 300 courses tagged by researchers along the same dimensions used in a previous study.

These datasets will form the basis for system evaluation and serve as the ground truth; for example, the expert-labeled XR course set is used to evaluate the filtering module's recall and precision.

3.3 Evaluation Protocol

The system is designed to be general-purpose and adaptable to curriculum analyses across various topics. For evaluation, we focus on XR as a case study. Accordingly, we configured the system to reflect the dimensions and category schemes used in the parallel study and specified a filtering condition: a course is considered relevant if its title or description contains any of the following keywords: "XR," "VR," "Virtual Reality," "Augmented Reality," "Extended Reality." We then ran the complete pipeline and evaluated performance across its three main components:

- *Extraction Accuracy:* Assesses the correct parsing of course title, description, university, department, and level from raw catalog data. Accuracy is measured by the proportion of correctly parsed entries, using string similarity thresholds.
- *Filtering Performance:* Compares system-identified relevant courses against an expert-defined ground truth using precision, recall, and F1-score (the harmonic mean of precision and recall).
- *Categorization Accuracy:* Evaluates how accurately the model classifies filtered courses into predefined categories for each dimension, based on a few-shot prompting strategy (i.e., the prompt supplies the model with task instructions plus a small number of illustrative examples rather than task-specific training data; Brown et al., 2020).

This evaluation ensures the system is both technically robust and practically aligned with academic needs in curriculum research and design.

4. SYSTEM ARCHITECTURE

The following system architecture was designed to be modular and user-oriented so that education researchers and administrators can adapt it to their needs. It is built on a Python backend that can execute a three-step pipeline: extraction, filtering, and categorization of courses. The following subsections describe each of these steps in detail, including their inputs, outputs, and transformations. Additionally, there is an independent evaluation module that computes performance metrics on the pipeline's output based on a user-provided set of ground-truth labels, according to the evaluation protocol described in the previous section. Finally, the system can be used via a Streamlit frontend which allows users to modify inputs, define configurations, and run any step of the pipeline independently (see Appendix A).

The main pipeline takes three key inputs:

- *A list of URLs* pointing to university course catalogs,
- *A list of keywords* to guide the filtering process,
- *A set of curricular dimensions* and their associated categories for the analysis.

The final output is an Excel file containing a detailed analysis of all courses that match the filtering criteria. Figure 2 illustrates the system architecture, while Appendix B summarizes the inputs, methods used, and the objectives of each component.

4.1 Extraction Module

The first step involves extracting course data from online university catalogs. The input to this module is a list of URLs, each pointing to a university's course catalog webpage. This is a web scraping task. Given the heterogeneity of web structures across institutions, building a deterministic script to accurately extract course information is non-trivial. We began by identifying repeated patterns. After analyzing the HTML of over 100 university catalogs – specifically those of institutions with AACSB-accredited business schools – we found that over two-thirds follow one of two common structures (screenshots in Appendix C):

- *Acalog Structure:* catalogs built using Modern Campus Catalog (formerly Acalog), Modern Campus's catalog management system. These consist of a structured list of course titles across multiple pages, organized by department. Each title links to a page with the course description.
- *Hierarchical Structure:* a single, alphabetically ordered page listing departments or categories, each linking to a page containing the full list of courses within that group. Courses are typically embedded in HTML "course blocks" containing both title and description.

Identifying these patterns allowed us to create deterministic parsers using BeautifulSoup and requests-html. For Acalog catalogs, the structure is highly regular and requires no heuristics. Hierarchical catalogs, however, vary more. To handle these, we adopted a key assumption: the course title is either explicitly marked in HTML or is the first element of a course block, with the description following. This assumption held true across all analyzed catalogs. As each catalog is scraped, course data is normalized and stored in a structured SQL table. Each row includes:

- Course Title
- Course Description
- University Name
- Department or School
- Academic Level (Undergrad or Grad)

This step generates the foundational course database for downstream analysis. It takes under 3 minutes per university, using a polite request rate to avoid overloading servers. Since university catalogs typically update annually, this step can be run infrequently. The remaining steps are much faster and can be run more regularly.

4.2 Filtering Module

The filtering module enables targeted analysis of course offerings based on user-defined areas of interest. It takes two main inputs: a filter specification and access to the previously created course database. The specification may consist of keywords (e.g., "virtual reality," "immersive technology") or full-text descriptions (e.g., "courses that introduce students to building XR environments or simulations"). To accommodate different user needs and improve relevance, the system supports both keyword-based and semantic filtering.

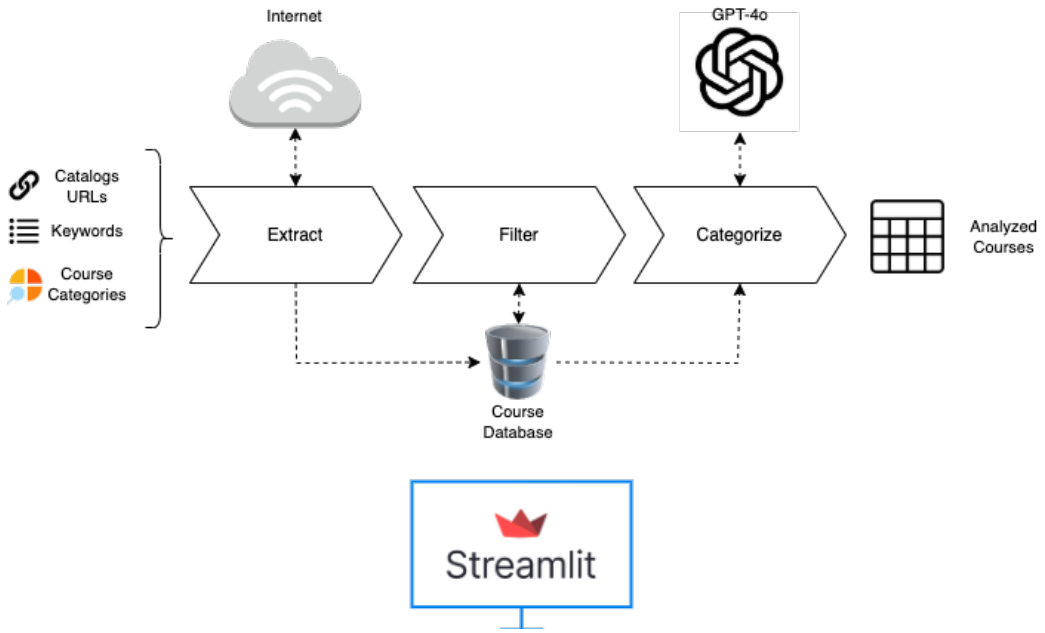


Figure 2. System Architecture

Keyword filtering uses TF-IDF-based scoring to identify high-signal matches between course descriptions and user queries. For semantic filtering, course descriptions and user-defined text can be embedded into vector representations using models such as Sentence-BERT or OpenAI's Ada. Cosine similarity is then used to measure semantic proximity and determine relevance. Filtering thresholds can be configured to prioritize either recall (for exploratory analysis) or precision (for highly targeted queries). The output of this stage is a reduced course table containing only the entries that match the user's interest profile, with confidence scores above a certain threshold.

4.3 Categorization Module

Once filtered, courses are passed into the categorization module, which enriches the data by classifying each course along custom curricular dimensions. These dimensions are defined by the user and can reflect conceptual frameworks such as "XR Use Case" (e.g., simulation, storytelling, prototyping), "Primary Discipline" (e.g., design, computer science, health sciences), or "Target Competency" (e.g., creative, technical, theoretical).

Each dimension is associated with a list of categories, described in natural language. The system uses few-shot prompting with a generative LLM (in this case GPT-4o) to classify each course. Prompt templates are populated with the course description and all category definitions, and the model is asked to select the best-fitting label. Following the few-shot prompting paradigm (Brown et al., 2020), each category definition is accompanied by brief illustrative examples (e.g., typical topics for each subject category, or a sample course for each course type), which anchor the model's interpretation of the categories without any fine-tuning. The iteratively refined template contains the course description together with the natural-language definitions of all categories for that dimension, enabling the model to apply nuanced, context-aware reasoning.

Categorization is performed independently for each dimension, allowing for multidimensional labeling. For example, a course may simultaneously be classified as "Technical" under the "Competency" dimension and "Prototyping" under the "Use Case" dimension. To ensure interpretability, the system can also request the model to output a rationale or explanation for its choice, which can be reviewed by researchers or institutional stakeholders.

4.4 Design Considerations

Several key principles guided the design of the system:

- *Modularity:* Each module can be run independently or in sequence, facilitating reuse across different research contexts or institutional settings.
- *Explainability:* Outputs at each stage include metadata and, where applicable, model-generated rationales to support human oversight and interpretation.
- *Scalability:* The system can process hundreds or thousands of course entries, making it suitable for large-scale cross-institutional studies.
- *Extensibility:* New dimensions, category definitions, or search interests can be integrated without code changes or retraining, leveraging the flexibility of prompt-based interaction with LLMs.

Taken together, these architectural components form a robust and extensible system for automated curriculum analysis. In the next section, we show how this approach aligns with previous work conducted in our companion study on XR in higher education. The full codebase and GenAI prompts for our system are available as open-source resources at https://github.com/luissacouto/curriculum-intelligence.

5. EVALUATION AND DISCUSSION

To assess the effectiveness of the proposed system, we conducted a structured evaluation focusing on precision, recall, and accuracy across each of the three core components: extraction, filtering, and categorization. Evaluation data were derived from our companion study on XR-related course offerings in AACSB-accredited institutions, where a manually curated dataset served as the ground truth for benchmarking. The extraction module demonstrated a high degree of accuracy, particularly when applied to the majority of standardized (Acalog-based and hierarchical) university websites. It should be noted that the study focused less on certain edge cases involving smaller private institutions employing customized JavaScript rendering. A targeted evaluation of 50 extracted courses confirmed the high fidelity of the process, with 100% accuracy in course titles and 98% in descriptions. In terms of speed, batch processing of institutional catalogs averaged under 3 minutes per university, indicating strong scalability for national or international studies.

In the filtering stage, the combination of keyword and semantic search yielded excellent results, with a precision of 100% and recall of 98.1% when benchmarked against the manually labeled XR course set. In contrast, the human-coded baseline – while also precise – demonstrated a considerably lower recall of just 36.3%, reflecting the limitations of manual review in exhaustively capturing relevant courses across heterogeneous catalogs. Of the 42 universities examined, the system identified more XR-relevant courses than the human coders in 71.4% of cases, matched them in 21.4%, and underperformed in only 7.1%. These results highlight the system's effectiveness not only in retrieving known courses but also in surfacing additional content that human reviewers missed. Notably, the hybrid approach – combining Sentence-BERT semantic scoring with keyword anchors – proved essential in capturing non-obvious courses that used less standard terminology.

The categorization module, powered by GPT-4o using few-shot prompts, accurately assigned courses to 10 different curricular dimensions. Evaluation against expert-coded categories showed an agreement rate above 88% across all dimensions, with higher consistency in more concrete categories and slightly lower scores in abstract ones. Importantly, the use of model-generated rationales enhanced transparency and interpretability, making the results more actionable for curriculum planners. A detailed comparison between human and AI topic labels across a manually annotated dataset is shown in the confusion matrix in Appendix D. It can be observed that the AI achieved very high agreement in key categories such as *Game Design* (96% with 50 exact matches) and *Communication & Media* (94%), confirming that the system is particularly effective for courses with clear disciplinary signals. Agreement was slightly lower in overlapping areas such as *XR Design & UX* (89%) and *XR Development* (81%). Overall, misclassifications tended to occur between semantically adjacent labels, rather than random mismatches – suggesting that the model's reasoning aligns closely with human judgments even when outcomes differ.

Similarly, when classifying course types – distinguishing between courses about XR, those that use XR devices in other areas, and those that do both – the system achieved an average agreement rate of 90% with the human annotations, based on the manually coded sample of 321 courses. The strongest alignment occurred for the "Both" category, where 115 out of 127 courses were correctly labeled. Most disagreement occurred around the "Uses XR" category, suggesting that ambiguity in course descriptions (e.g., where technology use is implied but not directly pointed out) may affect AI interpretation. The high agreement across both confusion matrices highlights the LLM's ability to capture nuanced curricular distinctions when guided by well-defined categories and few-shot examples.

Taken together, the system delivered fast, accurate, and scalable performance, validating its potential to replace manual content analysis in academic research and institutional decision-making. Its modularity and low-latency response times make it particularly suitable for iterative studies or dashboard applications that require up-to-date insights from dynamic course catalogs. While our evaluation focused on XR courses, the method can be readily applied to other emerging topics (cybersecurity, AI, etc.) by switching the filter criteria and category definitions. We anticipate a similar performance for well-defined topics, though additional testing in other domains and institutions (including non-U.S. or non-English catalogs) is a future need. An IS department chair could use this tool annually to benchmark their curriculum against those of peer institutions and identify if they are lagging in offering courses in areas such as XR or blockchain. Similarly, accreditation efforts (like those by AACSB or ABET) that require demonstrating curriculum currency could be informed by such automated analyses.

6. CONCLUSION

This work demonstrates a viable approach to automating curriculum analysis, which can greatly benefit educators and researchers by providing up-to-date, comprehensive views of how emerging technologies and other innovative topics are taught across institutions. Through our case study on XR-related academic programs, we show that our system delivers high levels of accuracy, consistency, and speed across its three core components. The system successfully parsed complex course catalog data, filtered content according to semantically defined criteria, and categorized offerings along custom curricular dimensions, all while significantly reducing the manual workload. Our approach can facilitate academic benchmarking (e.g., comparing XR offerings across universities), inform curriculum design decisions (identifying gaps or redundancies in a program), and even aid student advising by pinpointing where certain skills are taught. By building on the insights from a manually-coded study, we were able to validate the system against an established baseline, further demonstrating its reliability and value. This research highlights the transformative potential of generative AI in higher education analytics. As institutions navigate shorter innovation cycles and increasing pressure to align with industry needs, tools like the one presented can play a critical role in enabling evidence-based decision-making at scale. Methodologically, this work contributes to computing education research by introducing an AI-driven technique to gather and analyze curricular data, expanding the researchers' toolkit for evidence-based curriculum studies.

Looking ahead, future developments could include integrating interactive dashboards, expanding coverage to non-English institutions, and introducing feedback loops to improve categorization performance over time. Ultimately, this work lays the foundation for a new class of curriculum intelligence systems – combining the power of AI with the strategic needs of higher education.

7. REFERENCES

Awwad, E. (2024, February 27). *The Future of Higher Education Leadership Requires Innovation and Inclusion*. Forbes.com. https://www.forbes.com/councils/forbesbusinesscouncil/2024/02/27/the-future-of-higher-education-leadership-requires-innovation-and-inclusion

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). Language Models Are Few-Shot

Learners. In *Advances in Neural Information Processing Systems* (Vol. 33, pp. 1877-1901). Curran Associates.

Busulwa, R., Pathiranage, N. W., & Pickering, M. E. (2025). Integration of Critical Digital Technologies Into Hospitality and Tourism Curricula: Evidence From QS-Ranked Universities. *Journal of Hospitality and Tourism Technology*. https://doi.org/10.1108/JHTT-06-2024-0388

Ceccucci, W., Jones, K., Toskin, K., & Leonard, L. (2020). Undergraduate Business Analytics and the Overlap With Information Systems Programs. *Information Systems Education Journal*, 18(4), 22-32. https://doi.org/10.62273/NAVK3017

Chan, J., & Li, Y. (2025). Enhancing Higher Education With Generative AI: A Multimodal Approach for Personalised Learning. Proceedings of the 2025 International Conference on Advances in Education and Information Technology (pp. 50-57). https://doi.org/10.1007/978-981-96-8931-6_5

Dickey, E., & Bejarano, A. (2024). GAIDE: A Framework for Using Generative AI to Assist in Course Content Development. I*n 2024 IEEE Frontiers in Education Conference (FIE)* (pp. 1-9). IEEE. https://doi.org/10.1109/FIE61694.2024.10893132

Düdder, B., Fomin, V., Gürpinar, T., Henke, M., Iqbal, M., Janavičienė, V., Matulevičius, R., Straub, N., & Wu, H. (2021). Interdisciplinary Blockchain Education: Utilizing Blockchain Technology From Various Perspectives. *Frontiers in Blockchain*, 3(1), Article 578022. https://doi.org/10.3389/fbloc.2020.578022

Grimus, M. (2020). Emerging Technologies: Impacting Learning, Pedagogy and Curriculum Development. In S. Yu, M. Ally, & A. Tsinakos (Eds.), *Emerging Technologies and Pedagogies in the Curriculum* (pp. 127-151). Springer Singapore. https://doi.org/10.1007/978-981-15-0618-5_8

Gürpinar, T. (2025). Towards Web 4.0: Frameworks for Autonomous AI Agents and Decentralized Enterprise Coordination. *Frontiers in Blockchain*, 8, Article 1591907. https://doi.org/10.3389/fbloc.2025.1591907

Gürpinar, T., Verma, A., Leonard, L., Jones, K., & Ceccucci, W. (2025a). Blockchain Education: Evaluating Programs, Curricula, and Integration With Emerging Technologies in Higher Education. *Information Systems Education Journal*, 23(5), 20-29. https://doi.org/10.62273/WNRJ3445

Gürpinar, T., McKenna, A., Quadri, S., Jia, S. J., Sa-Couto, L., Lang, G. (2025b). Virtual, Augmented, and Extended Reality in Higher Education: Trends, Applications, Impacts. In *Proceedings of the 2025 Information Systems & Computing Academic Professionals (ISCAP) Conference* (Vol. 11, No. 6392).

Hoteit, L., El Hachem, M., Al-Hajjar, A., & Aouad, W. (2023, April 14). *Ten Lessons for Universities to Prepare Students for Long-Term Success*. BCG.com. https://www.bcg.com/publications/2023/ten-lessons-for-universities-prepare-students-future-of-work

Jovanović, J., Zamecnik, A., Barthakur, A., & Dawson, S. (2025). Curriculum Analytics: Exploring Assessment Objectives, Types, and Grades in a Study Program. *Education and Information Technologies*, 30(4), 4843-4866. https://doi.org/10.1007/s10639-024-13015-0

Naseer, F., Khalid, M. U., Ayub, N., Rasool, A., Abbas, T., & Afzal, M. W. (2024). Automated Assessment and Feedback in Higher Education Using Generative AI. In L. Tomei & M. B. Garcia (Eds.), *Global Innovations in Physical Education and Health* (pp. 433-461). IGI Global. https://doi.org/10.4018/979-8-3693-1351-0.ch021

Orellana, G., Orellana, M., Saquicela, V., Baculima, F., & Piedra, N. (2018). A Text Mining Methodology to Discover Syllabi Similarities Among Higher Education Institutions. In *2018 International Conference on Information Systems and Computer Science (INCISCOS)* (pp. 261-268). IEEE. https://doi.org/10.1109/INCISCOS.2018.00045

Zaki, N., Turaev, S., Shuaib, K., Krishnan, A., & Mohamed, E. (2023). Automating the Mapping of Course Learning Outcomes to Program Learning Outcomes Using Natural Language Processing for Accurate Educational Program Evaluation. *Education and Information Technologies*, 28(12), 16723-16742. https://doi.org/10.1007/s10639-023-11877-4

Zamecnik, A., Barthakur, A., Wang, H., & Dawson, S. (2024). Mapping Employable Skills in Higher Education Curriculum Using LLMs. In R. Ferreira Mello, N. Rummel, I. Jivet, G. Pishtari, & J. A. Ruipérez Valiente (Eds.), *Technology Enhanced Learning for Inclusive and Equitable Quality*

Education (Lecture Notes in Computer Science, Vol. 15160, pp. 18-32). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-72312-4_2

AUTHOR BIOGRAPHIES



Luis SaCouto is an Assistant Professor of Business Analytics & Information Systems at Quinnipiac University's School of Business. He holds a Ph.D. in Computer Science from the University of Lisbon, where his research on associative memory and vision earned the highest distinction. He has published over ten papers, authored a machine learning textbook, and taught AI and ML for six years as an invited lecturer. With experience at McKinsey and as Principal ML Scientist at DareData, his interests span both the foundational assumptions of learning systems and their applications in business and finance



Tan Gürpinar is an Assistant Professor of Business Analytics & Information Systems at Quinnipiac University's School of Business, with a focus on the impact of emerging technologies on organizations, people, and society. He serves as a Faculty Fellow in Residence and Advisor to the Living Learning Experience Leadership Council at Quinnipiac University. Dr. Gürpinar holds a Senior Researcher Position with the Fraunhofer Institute for Material Flow and Logistics, Dortmund, Germany, where technology research is transferred into practice to empower enterprise ecosystems. He also holds a Fellow Position at Yale University's Global Justice Program.



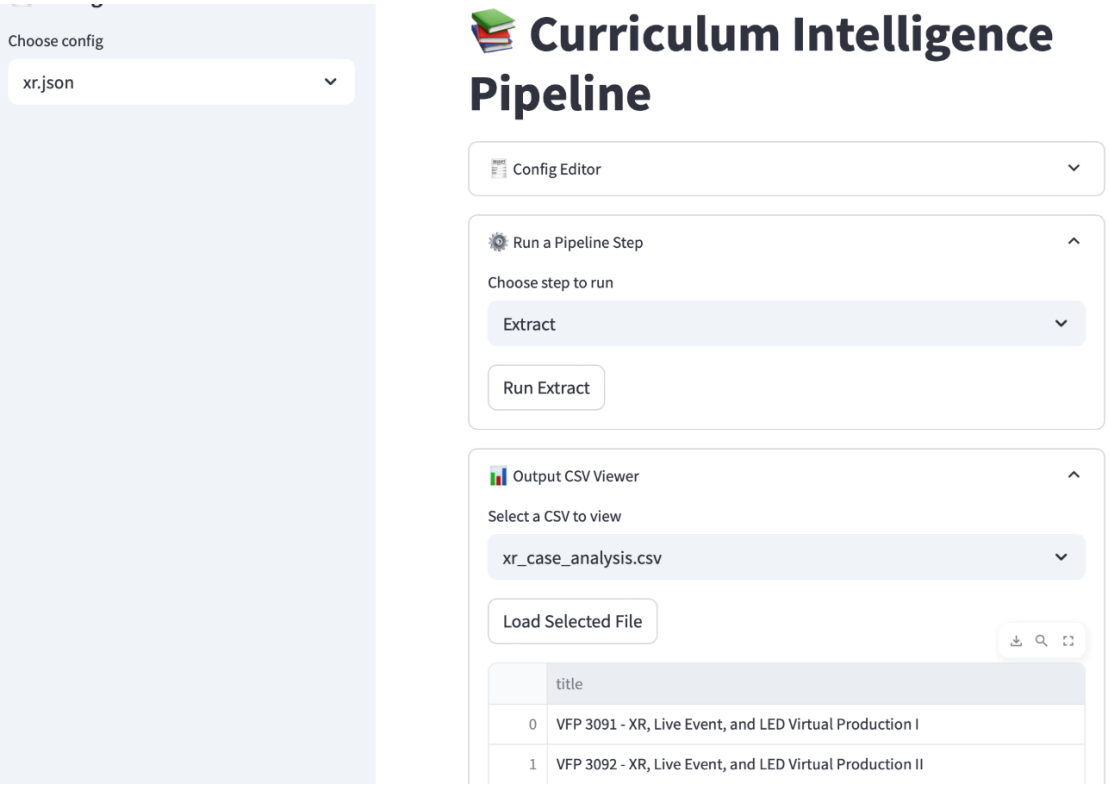
Shizhen (Jasper) Jia is an Assistant Professor of Business Analytics & Information Systems at Quinnipiac University's School of Business. His research interests are in human-AI interaction and ethical issues in innovation. He obtained his Ph.D. from the University of Florida and his Master of Science from Washington State University.



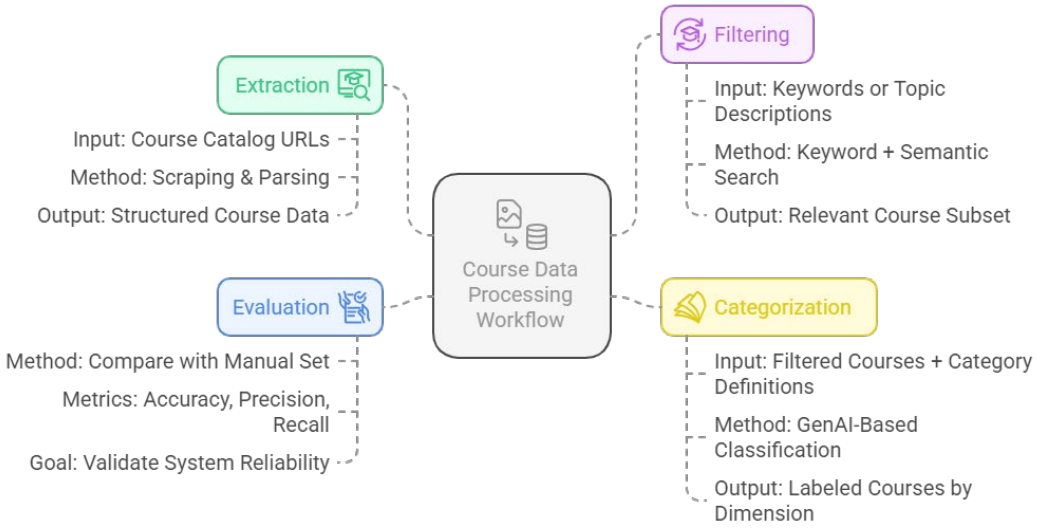
Sana Quadri is a student at Quinnipiac University majoring in International Business and Economics. Her Master is planned in Applied AI and Business Analytics. She is a Mentor at Quinnipiac's The Grove Learning Living Experience, supporting Big Idea Research Projects. Her research interests include emerging technology, culture, monetary policy, and development economics.

APPENDICES

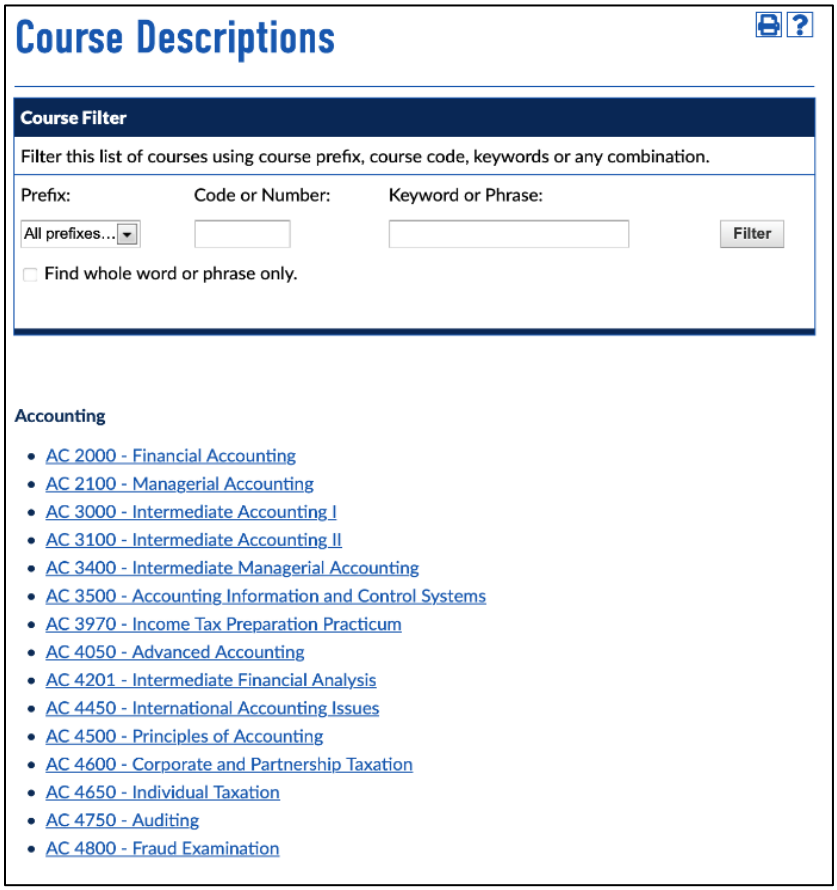
Appendix A. Screenshot of the Curriculum Intelligence System Frontend

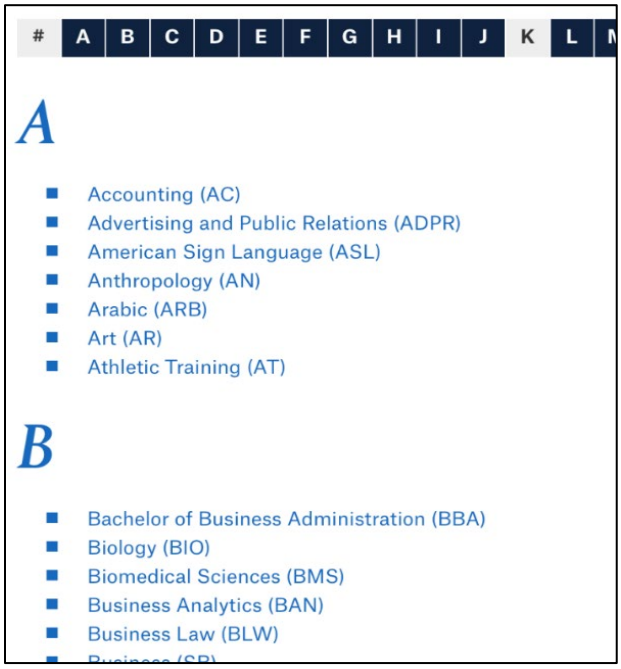


Appendix B. Input, Method, and Goal of the Curriculum Intelligence System



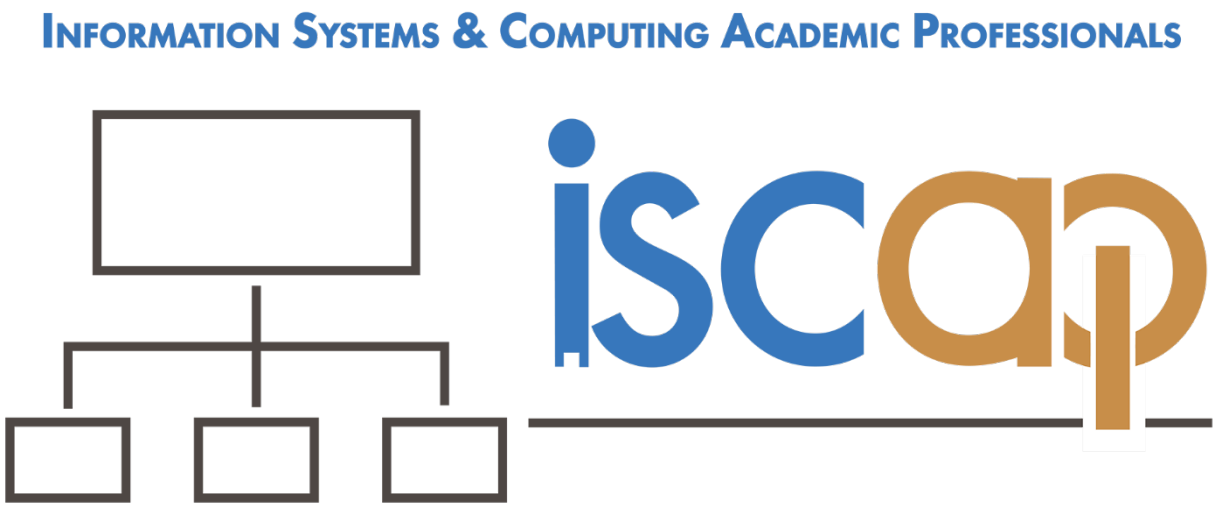
Appendix C. Acalog Catalog (Top) vs. Hierarchical Catalog (Bottom)





Appendix D: Confusion Matrices Displaying Human vs. AI Evaluations of XR Course Topics (Top) and Course Types (Bottom)





STATEMENT OF PEER REVIEW INTEGRITY

All papers published in the ***Journal of Information Systems Education*** have undergone rigorous peer review. This includes an initial editor screening and double-blind refereeing by three or more expert referees.

Copyright ©2026 by the Information Systems & Computing Academic Professionals, Inc. (ISCAP). Permission to make digital or hard copies of all or part of this journal for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial use. All copies must bear this notice and full citation. Permission from the Editor is required to post to servers, redistribute to lists, or utilize in a for-profit or commercial use. Permission requests should be sent to the Editor-in-Chief, ***Journal of Information Systems Education***, editor@jise.org.

ISSN: 2574-3872 (Online) 1055-3096 (Print)