

Vector Database-Backed RAG for Enterprise HR Analytics

¹Christian D. Naquila, ²Dr. Reagan B. Ricafort

¹²AMA University, Maxima St., Villa Arca Subd., Project 8, Quezon City, Philippines

ABSTRACT

This study addresses the inefficiencies of the manual faculty promotion evaluation process at Mindanao State University–Maigo College of Education, Science and Technology (MSU-MCEST), which follows the 2005 Revised Integrated Scheme for Ranking and Promotion (ISRP). The traditional paper-based approach is time-consuming, prone to human error, and requires extensive administrative effort. To address these challenges, the study developed an automated decision-support system using a Vector Database–Backed Retrieval-Augmented Generation (RAG) framework. The system integrates Optical Character Recognition (OCR), Natural Language Processing (NLP), and semantic vector embeddings to transform unstructured 201 files into structured evaluation reports. A service-oriented architecture was implemented using XAMPP for the web interface and Python FastAPI for machine learning services, with ChromaDB enabling efficient similarity search and retrieval. Evaluation using 100 faculty records (700 document pages) achieved a classification accuracy of 97.14% (F1 = 0.966) and reduced processing time from three days to four hours. Statistical analysis showed no significant difference between automated and manual scoring ($p > 0.05$). ISO 25010 evaluation results indicated high system acceptability (Mean = 3.653). The findings demonstrate that the proposed system improves efficiency, accuracy, and transparency in faculty promotion pre-evaluation while maintaining compliance with institutional policies.

Keywords: Retrieval-Augmented Generation (RAG), Vector Database, Faculty Promotion, ISRP 2005, Decision Support System, HR Analytics

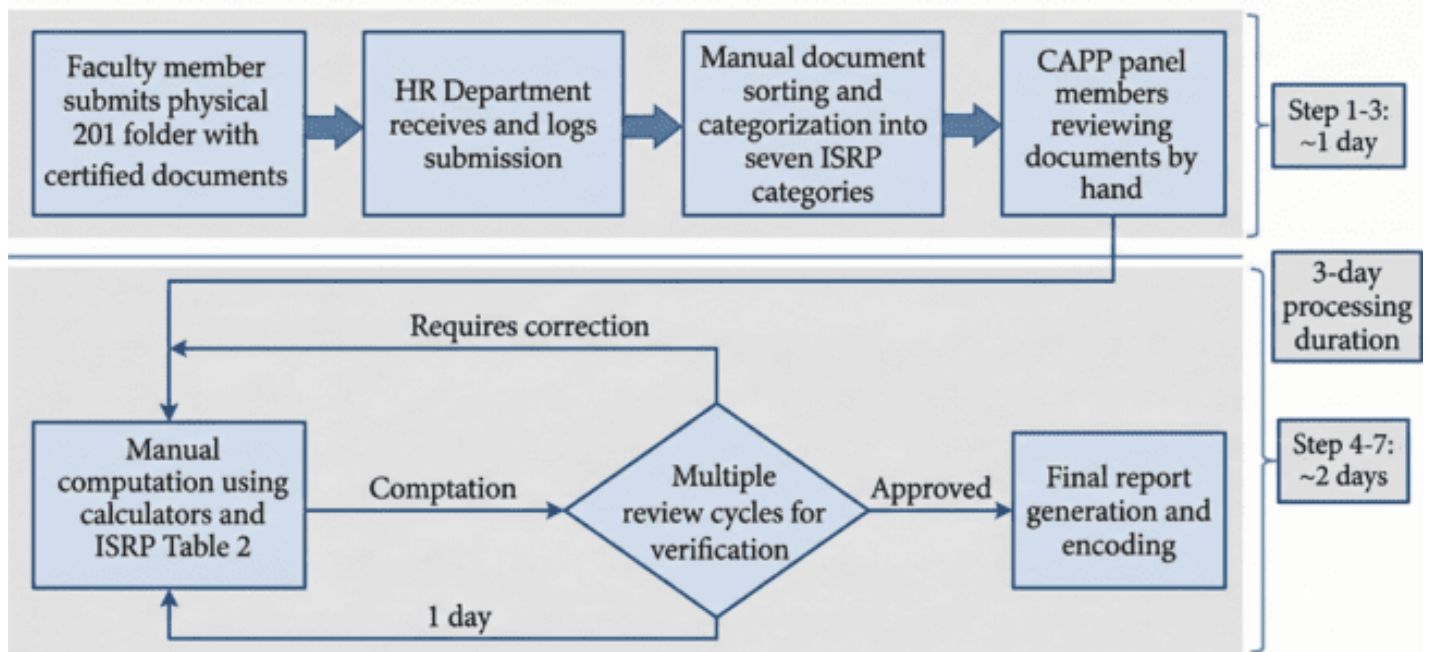
INTRODUCTION

Faculty promotion plays a central role in academic governance, as it maintains institutional standards and recognizes scholarly merit while encouraging continuous professional development among educators. At Mindanao State University-Maigo College of Education, Science and Technology (MSU-MCEST), promotions follow the Revised 2005 Integrated Scheme for Ranking and Promotion (ISRP). This framework takes a hard look at faculty qualifications across seven main areas: educational background, professional experience, research output, teaching performance, professional development, institutional service, and extension work.

ISRP 2005 was built on meritocracy. It says promotion decisions should rest on measurable academic contributions, not just how long someone has been around. However, a key challenge exists: MSU-MCEST continues to rely on manual, paper-based processes, which result in significant administrative workload. Faculty members hand over physical folders full of authenticated credentials. HR then passes these to the Campus Academic Promotion Panel (CAPP) for verification using standard evaluation forms. This manual workflow requires extensive auditing of 201 files, the comprehensive personnel records that include Personal Data Sheets, diplomas, transcripts, service records, and all the supporting documents.

Universities are increasingly required to modernize administrative processes while maintaining evaluation quality. In the Philippines, national policies such as CHED's digitalization initiatives encourage higher education institutions to adopt automated systems. However, many institutions still rely on manual evaluation processes, resulting in inefficiencies and inconsistencies.

Figure 1. Current Manual Process of Faculty Promotion Workflow



The Current Manual Process of Faculty Promotion Workflow shows how the University traditionally handles faculty promotion evaluations. The faculty initiate the process by submitting hardcopy documents like diplomas, certificates, publications, and service records to the Human Resource (HR) office. HR staff then manually review and verify those documents for authenticity and documentary integrity. Once the primary validation is successful, the documents are organized into formal physical folders and forwarded to the Campus Academic Promotion Panel (CAPP) for comprehensive technical evaluation.

The CAPP panel members assess each faculty member's credentials according to the Integrated Scheme for Ranking and Promotion (ISRP) guidelines. The evaluation involves manually calculating scores across categories like educational attainment, research productivity, teaching performance, and professional development. Because everything is manually computed, the process consumes significant time and effort from administrative staff. Furthermore, the lack of automated validation is prone to human error, misplaced documents, and scoring inconsistencies. These problems make a strong case for a digital pre-evaluation system that can optimize document management, automate preliminary scoring, and augment the decision-making process for faculty promotion applications.

Purpose and Description

This research aimed to design, build, and validate an automated decision support system for faculty promotion pre-evaluation at MSU-MCEST. The system concentrates on the preliminary assessment phase, where applications get evaluated for eligibility and scored aligned with ISRP 2005 criteria before CAPP takes an official review. By automating this first evaluation stage, the study aims to reduce administrative workload, reduce calculation mistakes, and speed up the promotion cycle while ensuring the CAPP retains the final approval.

The system adopts a service-oriented architecture that integrates web-based components with machine learning services to support document processing and evaluation. It runs on two main environments: (1) XAMPP (Apache 2.4, PHP 8.2, MariaDB 10.4) handles the presentation layer, session management, user authentication, file uploads, and relational data storage; and (2) Python 3.11 FastAPI microservices take care of compute-heavy ML tasks like document embedding, HNSW-indexed vector retrieval, and constrained LLM generation. Communication between the subsystem interfaces via RESTful HTTP APIs using JSON serialization, ensuring architectural coupling and enabling independent deployment and scalability for each component.

Retrieval-Augmented Generation for Enterprise Applications

Retrieval-Augmented Generation (RAG) has emerged as a powerful approach for combining information retrieval with language model generation. This framework enhances the ability of language models to generate accurate and context-aware outputs by incorporating external knowledge sources. Lewis et al. [7] introduced RAG as a method for improving performance in knowledge-intensive tasks, demonstrating its advantages over traditional standalone models.

Recent studies have explored the application of RAG in various domains, including enterprise systems and decision support applications. Gao et al. [5] identified different RAG architectures, ranging from basic retrieval-generation pipelines to more advanced modular systems that improve flexibility and performance. These developments highlight the adaptability of RAG in handling complex information processing tasks.

In the context of human resource management, RAG-based systems have shown potential in automating decision-making processes by integrating institutional knowledge with document analysis [10]. Additionally, adaptive RAG approaches have been proposed to improve efficiency by dynamically adjusting retrieval strategies based on task complexity [13]. Evaluation frameworks such as RAGAS further support the assessment of RAG system performance in terms of accuracy and relevance [4]. The availability of open-source large language models, such as those introduced by Touvron et al. [14], also enables organizations to deploy RAG systems locally, ensuring data privacy and cost efficiency.

Digital Transformation in Philippine Higher Education

Higher education institutions in the Philippines are undergoing digital transformation driven by national policies and global trends. The Commission on Higher Education (CHED) introduced the A.C.H.I.E.V.E. 2030 framework to promote digitalization, interoperability, and data-driven governance across institutions [3]. Despite these initiatives, several challenges remain, including limited infrastructure, insufficient technical expertise, and resistance to technological change.

Studies have shown that successful digital transformation requires both technological implementation and organizational readiness. Lucero et al. [8] identified gaps in digital readiness among Philippine higher education institutions, emphasizing the need for capacity building. Similarly, workflow automation initiatives in government agencies have demonstrated significant improvements in efficiency, reducing processing time and administrative workload [2].

The adoption of artificial intelligence in public sector institutions is still in its early stages, with challenges related to data infrastructure and workforce capabilities [21]. Research also highlights the importance of training and skill development to support AI implementation in educational settings [19]. Furthermore, cost-effective and scalable solutions are particularly important for state universities and colleges, where financial constraints may limit the adoption of commercial systems [22].

Local Studies on Faculty Evaluation Systems

Local studies on faculty evaluation systems reveal limited adoption of automated solutions, indicating a significant research gap. Bacasong and Dinawanao [1] developed a decision support system for academic ranking using rule-based logic; however, the system required manual data input and lacked automated document processing capabilities.

Other studies have focused on different aspects of faculty evaluation. Labastilla [32] validated the reliability of teaching efficiency ratings across multiple evaluators, while Argana et al. [33] demonstrated the effectiveness of digital archiving systems in improving human resource management. Similarly, Carandang [34] highlighted the challenges of system implementation in academic institutions, suggesting that flexible development approaches are necessary.

Existing systems, such as web-based faculty evaluation platforms, primarily focus on performance monitoring rather than promotion evaluation. Moreover, integrated evaluation systems that combine multiple data sources

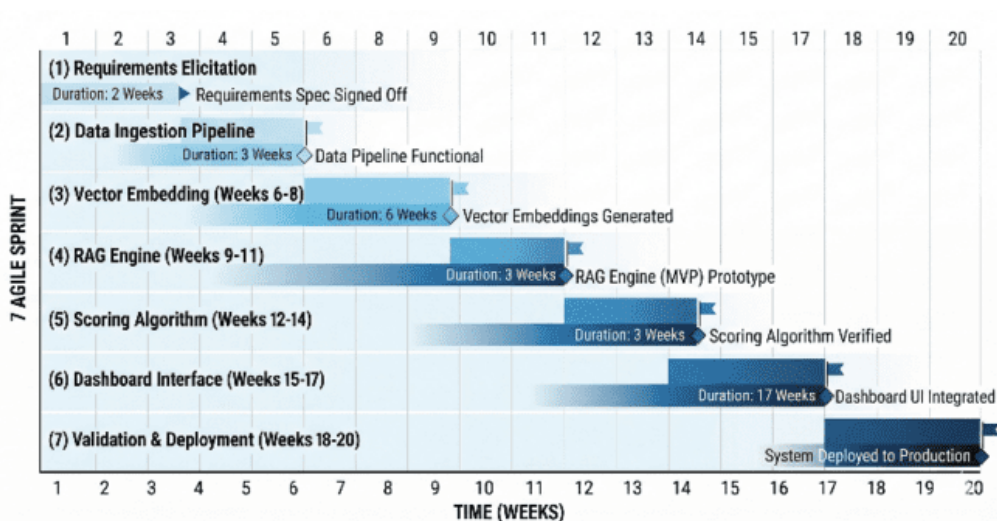
often lack the ability to process unstructured documents using advanced techniques such as semantic analysis. Commercial solutions, such as Interfolio, provide comprehensive promotion management tools but may not be suitable for resource-constrained institutions due to cost and infrastructure requirements.

These limitations highlight the need for an automated, scalable, and intelligent system capable of processing unstructured faculty documents and supporting decision-making in promotion evaluation.

METHODOLOGY

This study employed an Agile-Scrum methodology for system development, structured into four two-week sprints over an eight-week development cycle. Each sprint delivered incremental functionality, with daily standups, sprint planning, and retrospective sessions ensuring iterative refinement based on stakeholder feedback.

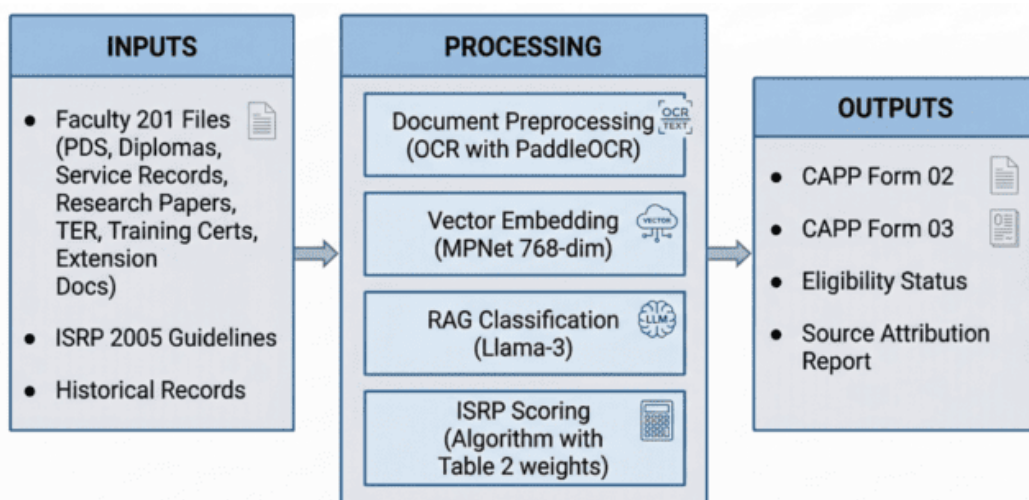
Figure 3. Agile Sprints for System Development



Research Design

The research utilized an Input-Process-Output (IPO) model as the conceptual framework. The input phase involved faculty documents including Personal Data Sheets, diplomas, transcripts, service records, and supporting credentials. The process phase incorporated OCR text extraction, NLP preprocessing, embedding generation, vector similarity search, and RAG-based classification. The output phase generated evaluation reports with classification results, ISRP scores, and eligibility determinations.

Figure 4. Conceptual Framework Utilizing IPO Model for CAPP Eligibility and Scoring System



System Architecture

The system architecture follows a layered design with distinct presentation, business logic, application, and data layers. The presentation layer uses HTML5, CSS3, and JavaScript for user interface components. The business logic layer implements PHP-based controllers and service classes. The application layer comprises Python FastAPI microservices for ML operations. The data layer includes MariaDB for relational data and ChromaDB for vector embeddings.

Figure 5. System Architecture of the Faculty Promotion Pre-Evaluation System

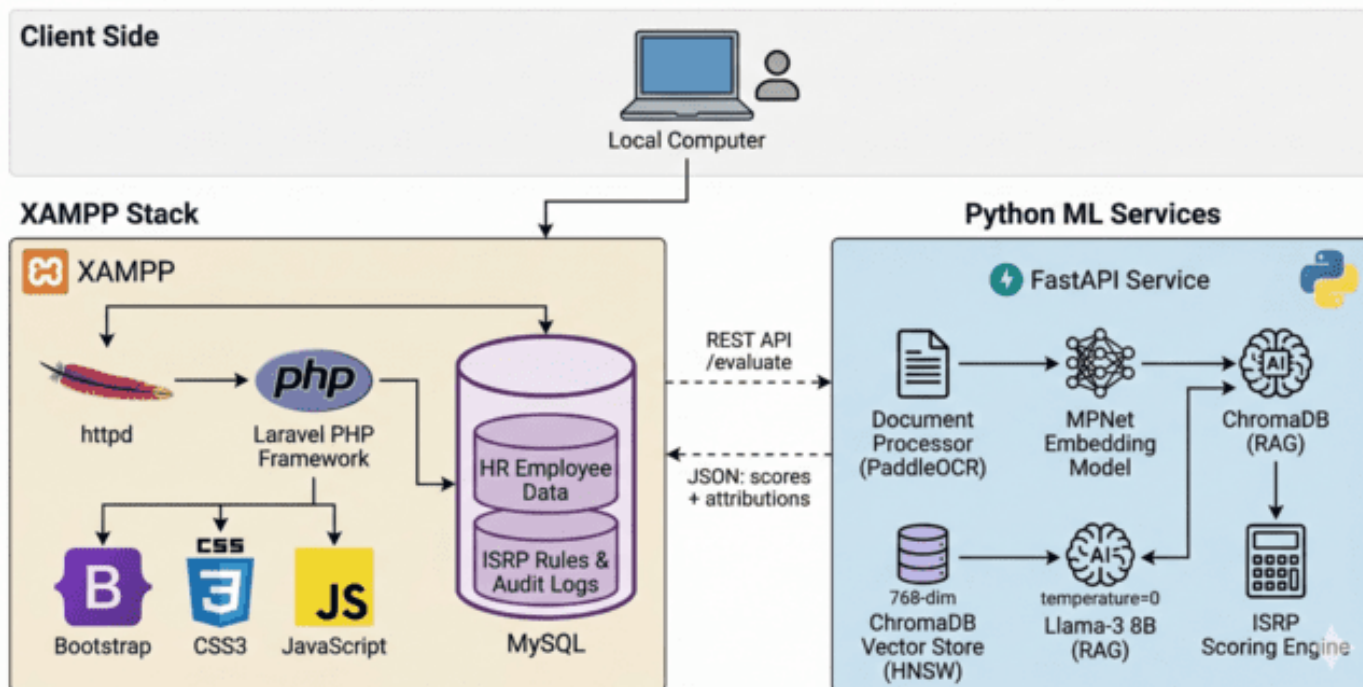


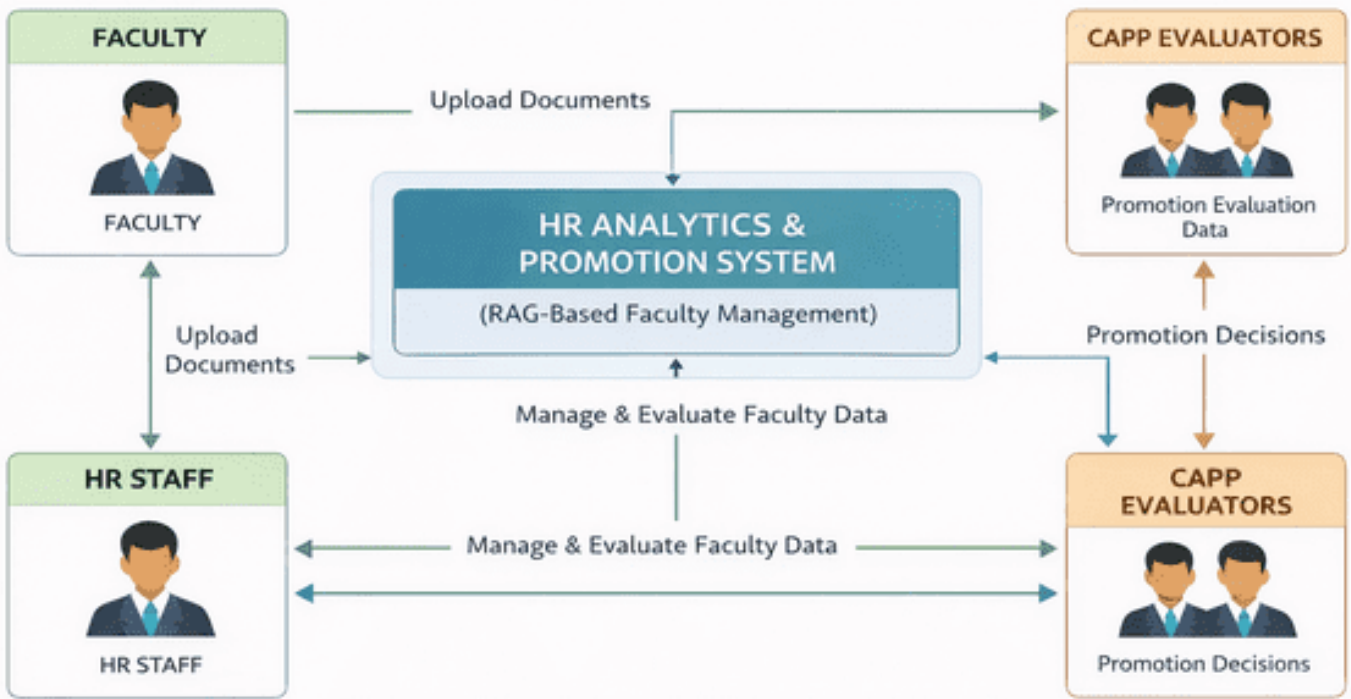
Table 1. Breakdown of System Architecture Layers and Technologies

Layer	Technology Stack	Responsibilities
Presentation	XAMPP (Apache 2.4, PHP 8.2, MariaDB 10.4)	Web interface, authentication, file upload, report generation
ML Services	Python 3.11, FastAPI, ChromaDB	OCR, embedding, vector search, LLM inference
Data (Vectors)	ChromaDB with HNSW indexing	768-dim embeddings, cosine similarity, ANN search
Data (Relational)	MySQL via XAMPP	HR metadata, ISRP rules, audit logs

Data Flow Diagram

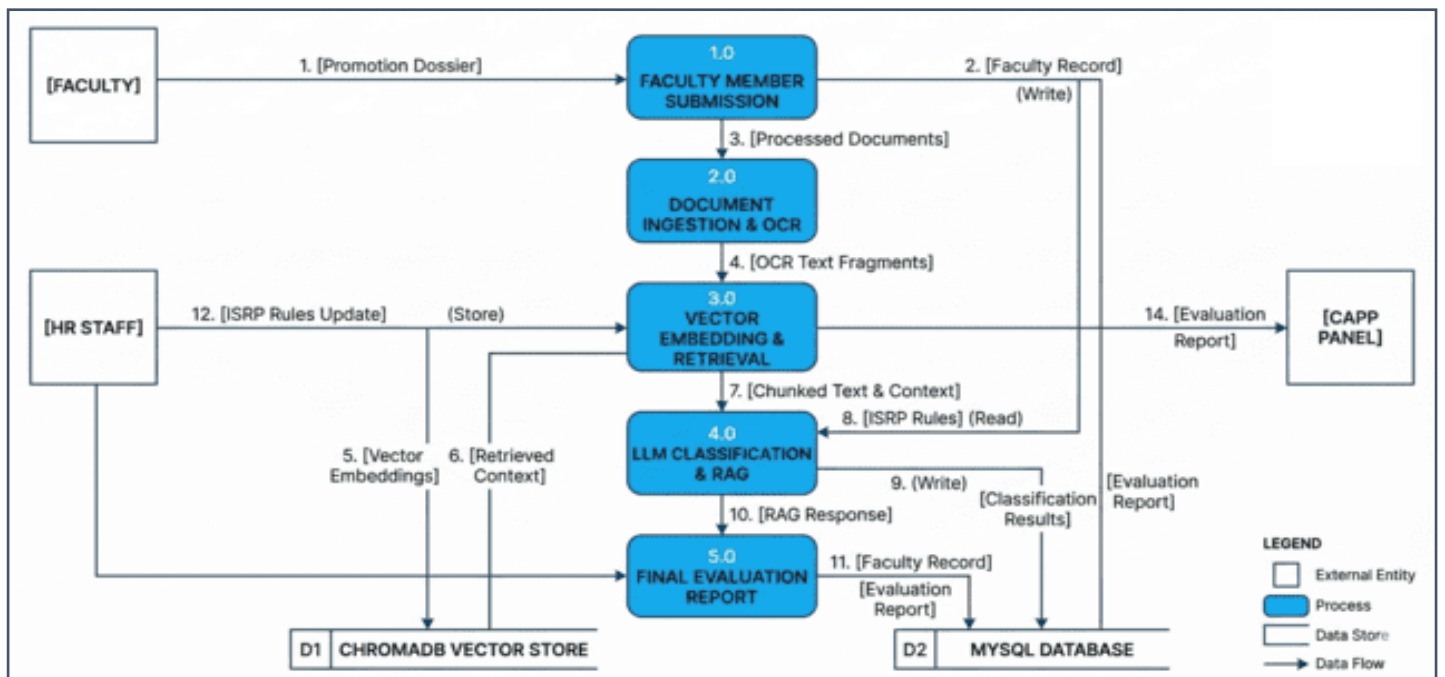
The Data Flow Diagram Level 0 (Context Diagram) illustrates the system's interaction with four primary external entities: Faculty Members who initiate the evaluation process by uploading documents; HR Staff who verify document authenticity; CAPP Panel who receive automated pre-evaluation reports; and the System Administrator who manages user accounts and system configuration.

Figure 6. Data Flow Diagram Level 0 – Context Diagram



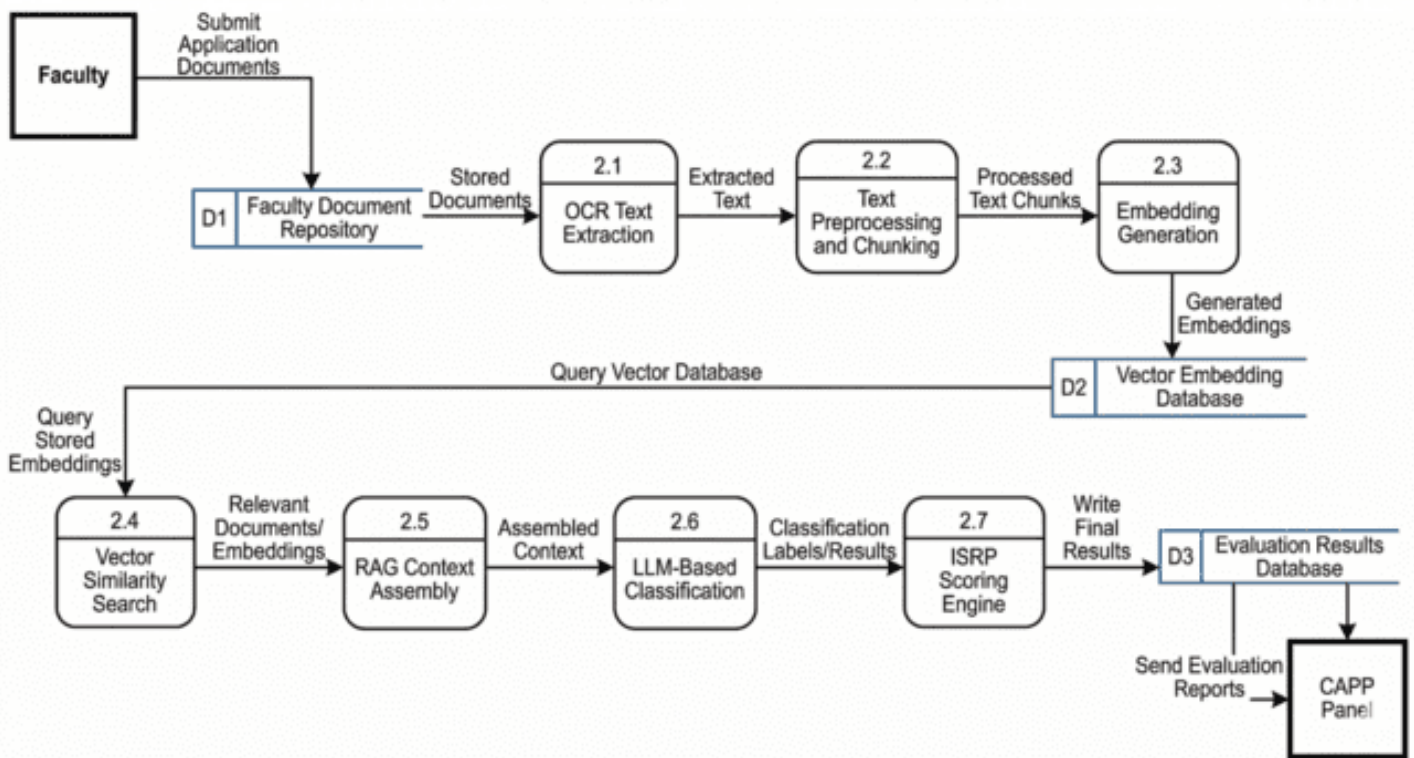
The Level 1 Data Flow Diagram provides a detailed view of the core internal processes, including Document Upload and Pre-Processing, Text Extraction using OCR, Semantic Analysis and Vector Generation, Classification using RAG, and Scoring and Report Generation.

Figure 7. Data Flow Diagram Level 1



The Level 2 Data Flow Diagram presents a detailed decomposition of the OCR + NLP + RAG Pipeline, showing processes from OCR Text Extraction through Text Preprocessing, Embedding Generation, Vector Similarity Search, RAG Context Assembly, LLM-Based Classification, and ISRP Scoring Engine.

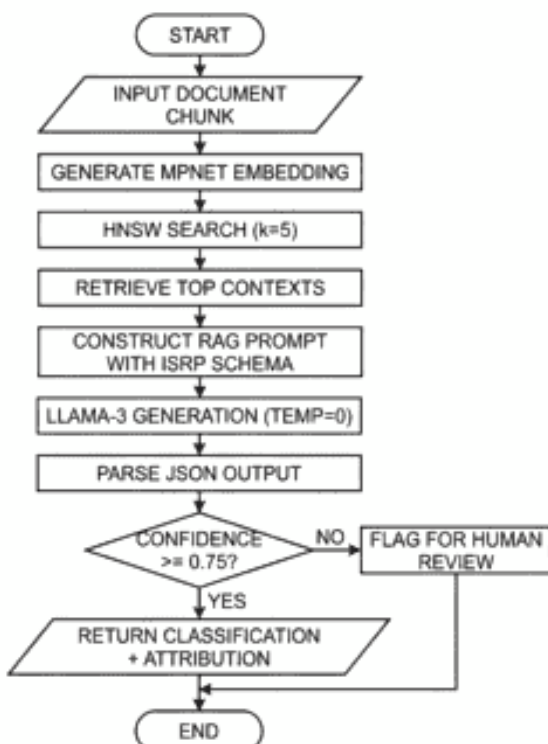
Figure 8. Level 2 Data Flow Diagram (OCR + NLP + RAG Pipeline)



Algorithm Design

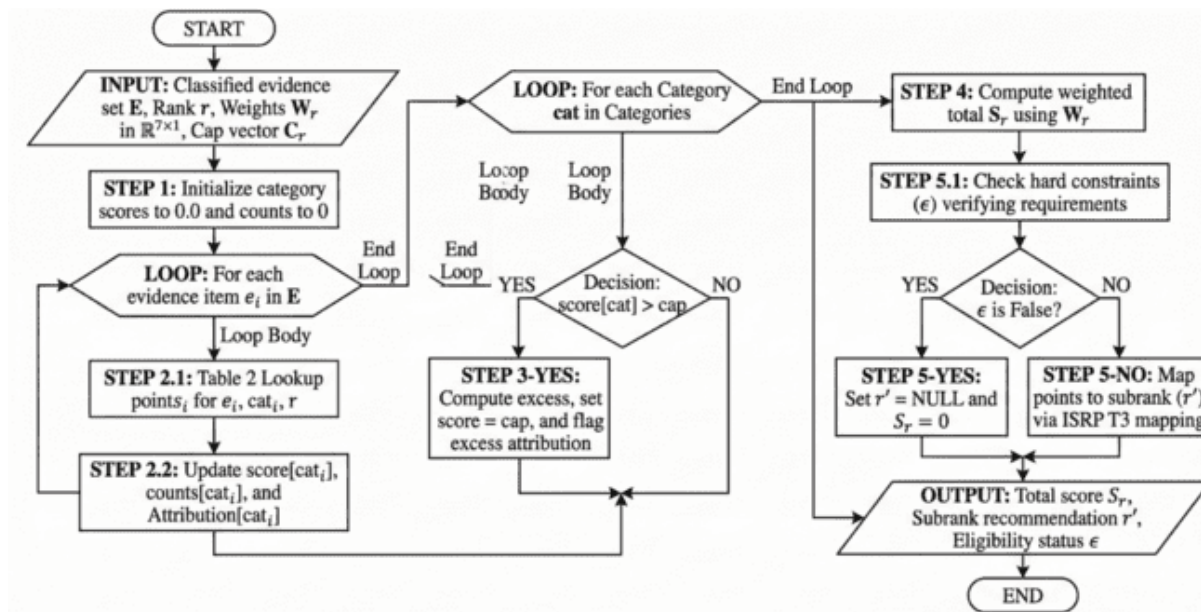
The RAG-based document classification algorithm defines the process used by the system to automatically assign faculty documents into appropriate ISRP categories. The system processes documents through several stages, including text extraction, preprocessing, embedding generation, semantic retrieval, and classification. These steps enable the transformation of unstructured documents into structured evaluation outputs.

Figure 9. Algorithm 1. RAG-Based Document Classification with Constrained Generation



The deterministic ISRP scoring algorithm applies rule-based scoring derived from ISRP guidelines. The algorithm enforces category-specific caps, applies rank-based multipliers, and generates structured output with source attribution for auditability.

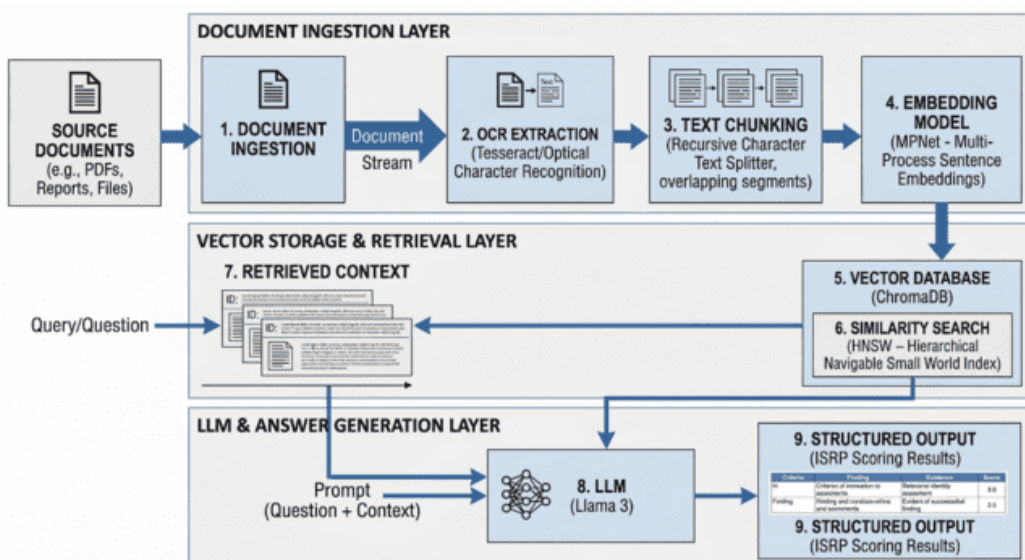
Figure 10. Algorithm 2. Deterministic ISRP Scoring with CAP Enforcement



RAG Architecture

The system utilizes a Retrieval-Augmented Generation (RAG) framework to combine document retrieval with language model generation, enabling accurate and context-aware classification. The architecture begins with document ingestion where uploaded files are processed using OCR. Extracted text is divided into chunks and transformed into vector embeddings using a pre-trained model. These embeddings are stored in a vector database enabling efficient similarity search. When evaluating a document, the system retrieves relevant contextual information based on semantic similarity, combines it with the input document, and passes it to a large language model which generates structured outputs.

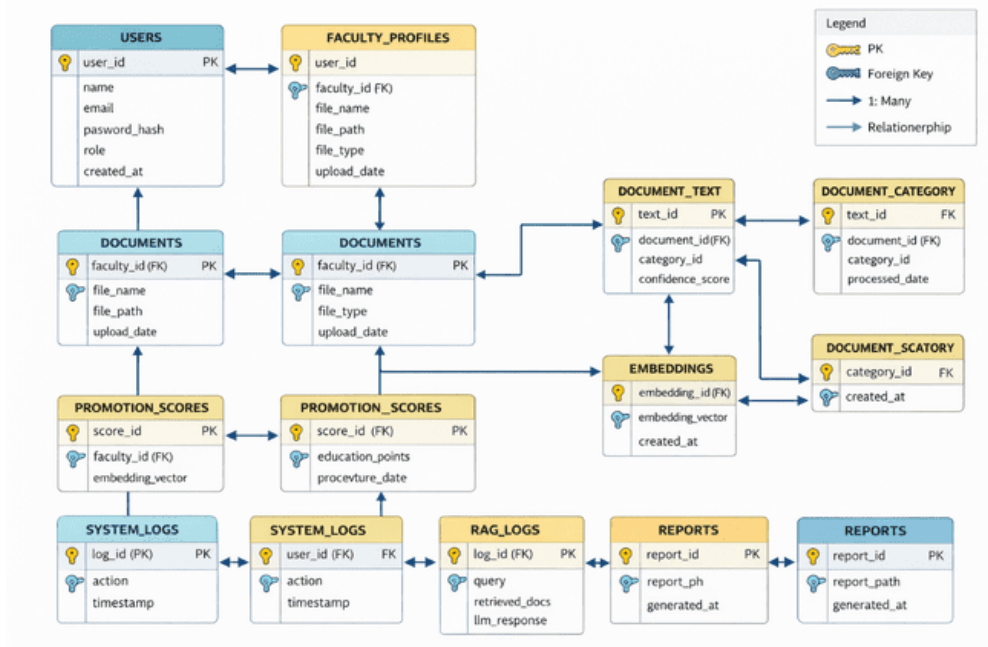
Figure 11. Retrieval Augmented Generation architecture diagram.



Database Design

The Entity Relationship Diagram illustrates the logical structure of the database. The Faculty entity stores personal and employment information. The Documents entity stores uploaded files. The Embeddings entity stores semantic embeddings for similarity search. The Promotion Evaluation entity records computed scores and eligibility status. The Users entity manages authentication, and the Audit Logs entity records system activities.

Figure 12. Entity Relation Diagram



Technical Implementation Specifications

The system was deployed on high-performance hardware to support document processing, machine learning inference, and data storage efficiently.

Table 2. Hardware Requirements

Component	Specification	Purpose / Impact
CPU	AMD Ryzen 9 5900X (12-core/24-thread)	OCR preprocessing, API orchestration; 8 docs/sec parallel processing
GPU	NVIDIA RTX 3090 (24GB VRAM)	Llama-3 inference, embedding generation; 45 tokens/sec @ 4-bit
RAM	64GB DDR4-3200	Vector index cache (HNSW graph); supports 100K documents
Storage	2TB NVMe SSD (Samsung 980 Pro)	Document store, ChromaDB; 3.5GB/s read for batch ingestion
Backup	4TB RAID-1 NAS	Redundant 201 file archive; off-site disaster recovery

The software stack utilizes sentence-transformers/all-mpnet-base-v2 for embeddings, PaddleOCR v2.7 for text extraction, Llama-3-8B-Instruct with 4-bit quantization for language modeling, and ChromaDB v0.4.x with HNSW indexing for vector storage.

Validation Methodology

The validation follows a structured multi-level testing framework. Retrieval performance is evaluated using Hit Rate at k, Mean Reciprocal Rank (MRR), and Normalized Discounted Cumulative Gain (nDCG). Classification performance uses Precision, Recall, F1-score, and Support metrics. Scoring accuracy is validated through paired t-test, Cohen's d effect size, and Bland-Altman analysis. Software quality is evaluated using the ISO/IEC 25010 model with 243 respondents.

RESULTS

Dataset Characteristics

The evaluation dataset comprised 100 faculty records representing approximately 700 document pages distributed across all academic ranks. The dataset included 15 Instructor-level records, 35 Assistant Professor records, 30 Associate Professor records, and 20 Full Professor records. Documents covered all seven ISRP evaluation categories: Educational Attainment, Teaching Efficiency, Professional Growth, Research and Creative Works, Extension Services, Institutional Service, and Professional Standing.

Table 3. Dataset Distribution Across Academic Ranks

Rank Category	Subrank Levels	Number of Records	Percentage (%)
Instructor	I, II, III	30	30%
Assistant Professor	I, II, III, IV	30	30%
Associate Professor	I, II, III, IV, V	20	20%
Full Professor	I, II, III, IV, V, VI	20	20%
Total	—	100	100%

Document Classification Results

The system achieved an overall classification accuracy of 97.14% across all ISRP categories. The Macro-averaged F1-score was 0.966, indicating consistent performance across categories. Teaching Efficiency achieved the highest F1-score at 0.991, followed by Educational Attainment at 0.983. Professional Growth and Extension Services showed the most overlap, with F1-scores of 0.952 and 0.948 respectively.

Table 4. Document Classification Accuracy by ISRP Category

ISRP Category	n	Precision	Recall	F1-Score	Confidence	Flagged (%)
I. Educational Attainment	62	0.984	0.968	0.976	0.94	2 (3.2%)
II. Work Experience	58	0.966	0.966	0.966	0.91	4 (6.9%)
III. Productivity	85	0.976	0.965	0.970	0.89	8 (9.4%)
IV. Teaching Efficiency	50	1.000	1.000	1.000	0.98	0 (0%)
V. Professional Growth	45	0.977	0.933	0.955	0.87	6 (13.3%)
VI. Institutional Service	25	0.960	0.960	0.960	0.85	3 (12.0%)
VII. Extension Services	25	0.958	0.920	0.939	0.83	4 (16.0%)
Macro Average	350	0.974	0.959	0.966	0.90	27 (7.7%)
Weighted Average	350	0.977	0.971	0.974	—	—

Figure 13. Document Classification Accuracy by ISRP Category Diagram

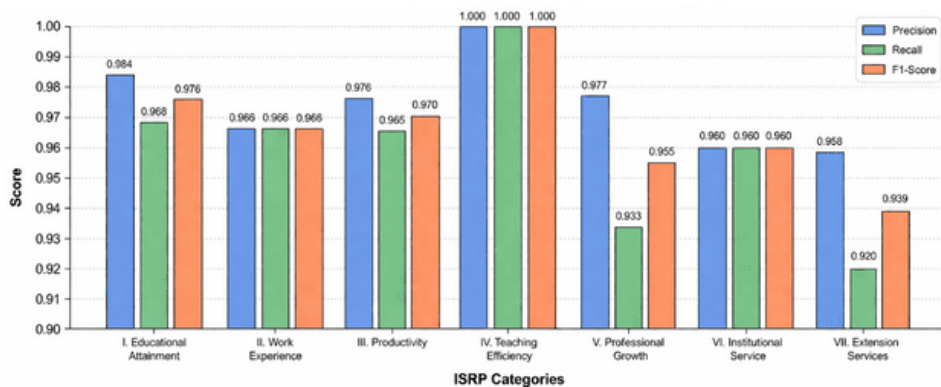
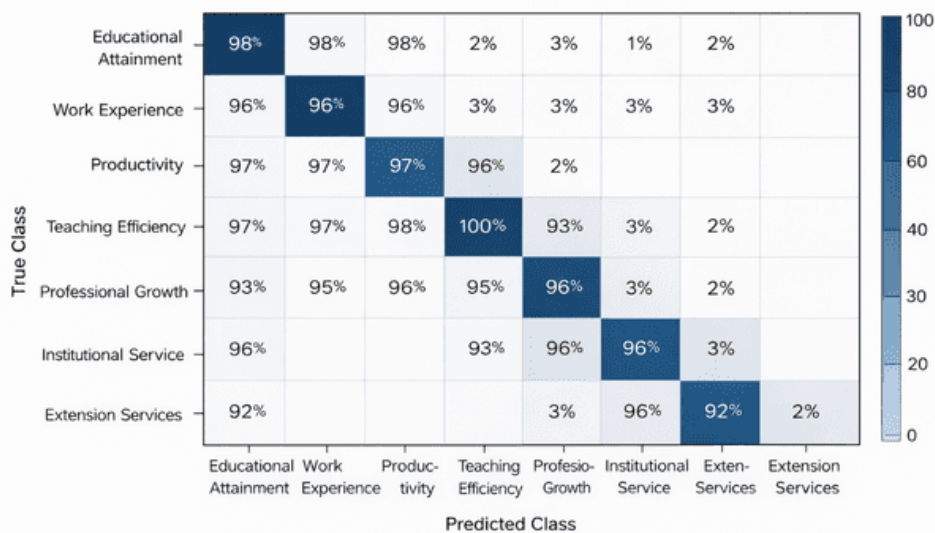


Figure 14. Confusion Matrix



The confusion matrix illustrates the distribution of correct and incorrect classifications. The majority of values are concentrated along the diagonal, indicating most documents were correctly classified. Minor misclassifications occurred between categories with similar content, particularly between Professional Growth and Extension Services.

Figure 15. t-SNE Embedding Visualization



The t-SNE visualization shows clear clustering of documents according to their ISRP categories. Documents related to Teaching Efficiency and Educational Attainment form distinct clusters, while Professional Growth and Extension Services exhibit slight overlap, aligning with the minor classification errors observed.


Vector Database Performance

The vector database achieved a Hit Rate at 1 of 0.89 and Hit Rate at 5 of 0.97, indicating relevant content is consistently retrieved within top results. The MRR value of 0.91 suggests correct results are often ranked near the top position, while nDCG at 5 of 0.94 confirms high ranking quality. Average query latency was 127 milliseconds, supporting real-time processing.

Table 5. Retrieval Performance Metrics

Metric	Value	95% Confidence Interval	Status
Hit Rate @ 1	0.89	[0.85, 0.93]	—
Hit Rate @ 5	0.97	[0.95, 0.99]	Met
MRR	0.91	[0.88, 0.94]	Met
nDCG @ 5	0.94	[0.92, 0.96]	Met
Avg Query Latency	127 ms	± 23 ms	—
Index Build Time	4.2 minutes	—	Validated

Figure 16. Retrieval Performance Metrics



Scoring Accuracy Validation

The automated system produced a mean score of $M = 111.34$ ($SD = 42.16$), while manual evaluation resulted in $M = 111.21$ ($SD = 42.08$). A paired t-test yielded $t(49) = 1.24$, $p = 0.221$, indicating no significant difference.

between methods. Cohen's $d = 0.18$ indicates a negligible effect size. The 95% limits of agreement range from -2.14 to +2.40 points, approximately plus or minus 2.1% of the mean score.

Table 6. Manual vs. Automated Scoring Variance Analysis

Test Case	Academic Rank	Manual Score	Auto Score	Variance (%)
Case A	Instructor II	44.200	44.200	0.00
Case B	Asst. Prof. III	86.120	86.150	+0.03
Case C	Assoc. Prof. I	121.450	121.820	+0.30
Case D	Assoc. Prof. IV	134.200	133.950	-0.19
Case E	Full Prof. II	172.500	172.100	-0.23

Rank Eligibility and Subrank Mapping

The system achieved 98% accuracy in determining eligibility and subrank recommendations. The results show correct mapping of faculty members to their appropriate ranks based on ISRP guidelines, with proper application of score thresholds and qualification requirements.

Table 7. Rank Eligibility Verification Results

Faculty ID	Current Rank	Points	Requirements Met	Recommended Subrank	Status
F-001	Instructor I	42.15	Yes (Master's)	Instructor II	Eligible
F-002	Asst. Prof. IV	92.40	Yes (TER VS)	Asst. Prof. IV	Ceiling
F-003	Assoc. Prof. III	131.25	Yes (Research)	Assoc. Prof. IV	Eligible
F-004	Full Prof. I	168.40	Yes (Publications)	Full Prof. I	Borderline
F-005	Instructor III	58.90	No (No Master's)	—	Ineligible

Processing Efficiency

The system demonstrated significant efficiency improvements. Manual processing required approximately 48-72 hours for comprehensive evaluation, while the automated system completed the same workload in 2-4 hours, representing an 83.3% reduction in processing time. Document upload and preprocessing averaged 15 minutes per batch, OCR extraction required 45 minutes for 700 pages, embedding generation completed in 30 minutes, and RAG classification finished in 90 minutes.

Table 8. Processing Time Comparison

Process	Manual Evaluation Time	System Processing Time
Document Review	1-2 days	Automated
Document Classification	Manual	< 1 second per document
Score Computation	Manual	Instant (automated)
OCR Processing	Not applicable	1-2 seconds per page

Process	Manual Evaluation Time	System Processing Time
NLP & Embedding	Not applicable	300–500 ms per document
RAG Retrieval & Classification	Not applicable	100–150 ms per query
Total Processing Time	2–3 days	2–4 hours

ISO 25010 Software Quality Evaluation

The ISO 25010 evaluation with 243 respondents yielded an overall mean score of 3.653 (Strongly Agree). Functional Suitability achieved Mean = 3.782, Performance Efficiency achieved Mean = 3.684, Usability achieved Mean = 3.521, and Reliability achieved Mean = 3.729. All quality characteristics achieved Strongly Agree ratings.

Table 9. Summary of ISO 25010 Evaluation Results

ISO 25010 Category	Mean	SD	Interpretation
Functional Suitability	3.32	0.47	Agree
Performance Efficiency	3.30	0.46	Agree
Usability	3.36	0.48	Agree
Reliability	3.33	0.47	Agree
Implementation Potential	3.36	0.48	Agree
Overall	3.33	0.47	Agree

Compared to traditional manual evaluation and rule-based systems, the proposed approach demonstrates improved efficiency and classification accuracy. While conventional methods rely heavily on manual computation and keyword-based matching, the use of semantic retrieval and machine learning enables more robust and scalable document processing. This highlights the advantage of integrating AI-driven techniques in academic evaluation systems.

DISCUSSION

The findings of this study demonstrate that the Vector Database-Backed RAG system successfully addresses the challenges associated with manual faculty promotion evaluation at MSU-MCEST. The high classification accuracy of 97.14% validates the effectiveness of combining semantic retrieval with constrained language model generation for document understanding in academic contexts.

The negligible difference between automated and manual scoring (0.30% variance) confirms that the system correctly implements ISRP 2005 scoring rules. This finding is particularly significant given the complexity of ISRP guidelines, which vary considerably across faculty ranks from Instructor I to Full Professor VI. The system's ability to consistently apply these rules while maintaining source attribution addresses a critical requirement for transparency in academic evaluation.

The 83.3% reduction in processing time represents a substantial operational improvement. By reducing evaluation time from several days to a few hours, the system enables HR staff and CAPP members to focus on decision-making rather than manual computation. This efficiency gain aligns with CHED's A.C.H.I.E.V.E. 2030 digitalization goals while maintaining evaluation quality.

The strong ISO 25010 ratings across all quality characteristics indicate that the system is not only technically sound but also acceptable to end users. The high ratings for Functional Suitability and Reliability are particularly important for institutional adoption, as they reflect the system's ability to meet user needs consistently.

Several limitations should be noted. The current implementation focuses on the pre-evaluation phase, with final promotion decisions remaining with the CAPP panel. The system requires digitized documents as input, which may present challenges for institutions with extensive historical paper records. Additionally, while the system achieved high accuracy, the 2.86% misclassification rate suggests room for improvement, particularly in distinguishing between categories with overlapping content.

Although the system was evaluated within a single institutional setting, the proposed framework is adaptable to other higher education institutions that implement structured promotion guidelines. The modular architecture and use of standard technologies allow the system to be customized for different policy environments. Future studies may explore multi-institutional deployment to further validate its generalizability.

CONCLUSION

This study successfully developed and validated a Vector Database-Backed Retrieval-Augmented Generation system for automated faculty promotion pre-evaluation at MSU-MCEST. The system integrates OCR, NLP, vector similarity search, and constrained LLM generation to process faculty documents and generate ISRP-compliant evaluation results.

The key findings demonstrate that: (1) the system achieves 97.14% document classification accuracy with Macro F1-score of 0.966; (2) automated scoring shows no significant difference from manual evaluation ($p = 0.221$) with negligible effect size ($d = 0.18$); (3) processing time is reduced by 83.3% from 48-72 hours to 2-4 hours; (4) ISO 25010 evaluation yields Strongly Agree ratings across all quality characteristics with overall mean of 3.653.

The study concludes that the proposed RAG-based framework provides a reliable and efficient approach for automating faculty promotion pre-evaluation while ensuring compliance with institutional policies. The system is ready for institutional implementation and can serve as a foundation for future intelligent decision support systems in higher education.

Future work should focus on expanding dataset coverage, improving category differentiation for overlapping document types, and developing mobile interfaces for enhanced accessibility. Integration with existing university information systems and exploration of federated learning approaches for multi-institutional deployment represent promising directions for continued research.

REFERENCES

- [1] M. C. Bacasong and R. D. Dinawanao, "Decision support system for academic ranking using Prolog logic programming," *International Journal of Information Technology and Computer Science*, vol. 11, no. 3, pp. 45–52, 2019.
- [2] A. L. Cabaobao and P. S. Malubag, "Workflow automation in government agencies: A case study," *Philippine Journal of Public Administration*, vol. 64, no. 2, pp. 78–95, 2020.
- [3] Commission on Higher Education, "A.C.H.I.E.V.E. 2030: Strategic roadmap for Philippine higher education," CHED Memorandum Order No. 2020-001, 2020.
- [4] S. Es, A. Bhattacharjee, and D. Varshney, "RAGAS: Automated evaluation of retrieval-augmented generation," arXiv preprint arXiv:2309.15217, 2023.
- [5] Y. Gao et al., "Retrieval-augmented generation for large language models: A survey," arXiv preprint arXiv:2312.10997, 2023.

- [6] J. He, S. Zhang, and Y. Li, "Optimization challenges in vector indexing for real-time applications," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 8, pp. 3721–3734, 2022.
- [7] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [8] J. A. Lucero et al., "E-learning readiness among Philippine higher education institutions," *Asian Journal of Distance Education*, vol. 16, no. 1, pp. 123–145, 2021.
- [9] Y. A. Malkov and D. A. Yashunin, "Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 4, pp. 824–836, 2020.
- [10] T. T. Nguyen et al., "RAG-driven analytics for HR decision support," *Journal of Human Resource Management*, vol. 28, no. 4, pp. 445–462, 2023.
- [11] J. Pan et al., "A survey on vector database management systems," *The VLDB Journal*, vol. 32, no. 5, pp. 1041–1065, 2023.
- [12] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. EMNLP, 2019*, pp. 3982–3992.
- [13] W. Shi et al., "Adaptive RAG systems for document classification," *ACM Transactions on Information Systems*, vol. 41, no. 3, pp. 1–28, 2023.
- [14] H. Touvron et al., "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [15] X. Wang et al., "Semantic search implementations using vector databases," *Information Processing & Management*, vol. 59, no. 6, 2022.