



Concept of an AI-assisted Regulations and Student Support Chatbot for the Faculty of Engineering at the University of Mauritius

Dr. Yogesh Beeharry* , Department of Electrical and Electronic Engineering, Faculty of Engineering, University of Mauritius, Réduit, Mauritius (y.beeharry@uom.ac.mu)

Dr. A. K. Ragen , Department of Chemical and Environmental Engineering, Faculty of Engineering, University of Mauritius, Réduit, Mauritius (ak.ragen@uom.ac.mu)

Dr. A. Khoodaruth , Department of Mechanical and Production Engineering, Faculty of Engineering, University of Mauritius, Réduit, Mauritius (a.khoodaruth@uom.ac.mu)

*Corresponding author

Abstract

The growing complexity of programme regulations, administrative procedures, and academic requirements within higher education institutions often leaves students struggling to access accurate and timely information. This paper presents the concept of an artificial intelligence (AI)-assisted Regulations and Student Support Chatbot designed specifically for the Faculty of Engineering at the University of Mauritius. The proposed system aims to serve as an accessible digital assistant capable of responding to student queries related to programme regulations, progression rules, assessment policies, submission guidelines, and general administrative processes. Leveraging advances in large language models (LLMs) and retrieval-augmented generation (RAG), the chatbot would integrate faculty-specific documents—such as student handbooks, programme regulations, and official University policies—into a searchable knowledge base to deliver precise and contextually grounded answers.

Beyond providing information, the chatbot is conceptualised as a tool to enhance student success by reducing uncertainty, improving regulatory literacy, and supporting self-directed learning. It aims to complement rather than replace the academic and administrative support structures by streamlining routine interactions. This would allow staff members to focus on higher-value tasks. The paper outlines the system architecture, data preparation workflow, ethical considerations, and proposed deployment strategy. Emphasis is laid on transparency, accuracy, and responsible use.

By addressing information bottlenecks and improving communication between students and the institution, the proposed AI-assisted chatbot will contribute to a more supportive and efficient learning environment. It aims to illustrate how tailored AI solutions can strengthen student engagement, enhance operational efficiency, and support institutional goals in a resource-constrained higher education context.

Keywords: AI-Assisted; Chatbot; LLMs; RAG; Student Support

1. Introduction

Higher education institutions across the world are undergoing rapid transformation as the integration of artificial intelligence (AI) reshapes administrative processes, learning environments, and student support systems. Universities face growing demands to improve student experience, streamline information delivery, and increase operational efficiency while maintaining high academic standards. There is a need for clear, accessible, and reliable student guidance for engineering faculties where programme regulations, progression rules, and assessment policies are often complex and frequently updated. Students frequently struggle to interpret regulatory documents or navigate institutional procedures, creating bottlenecks that impact academic performance,

administrative workload, and overall student satisfaction (Andrenucci & Sneiders, 2005).

At the University of Mauritius (UoM), the Faculty of Engineering (FoE) offers multiple BEng and postgraduate programmes with distinct rules governing credit accumulation, prerequisites, resits, retake exams, assessments, work placement requirements, and degree classifications. These details are typically distributed across programme regulations, handbooks, departmental guidelines, Senate regulations, and faculty notices. Staff members provide essential support, but increasing enrolment numbers and resource constraints limit the capacity for personalised assistance. Students routinely report confusion around progression criteria, module choices, project regulations, industrial training requirements and

assessment timelines, which are issues that contribute to avoidable academic risks. As in many institutions, a significant portion of queries received by administrative and academic staff are repetitive, procedural, and document-based (Winkler & Soellner, 2018).

Recent advances in natural language processing (NLP) and large language models (LLMs) have introduced new opportunities for educational institutions to automate routine, information-based student services while improving accuracy and accessibility. AI-driven chatbots have demonstrated strong potential in handling high-volume queries, supporting self-regulated learning, and enhancing student engagement through immediate, interactive feedback (Holmes et al., 2019). Moreover, retrieval-augmented generation (RAG) systems enable AI models to provide grounded, institution-specific answers by retrieving content from official documents before generating a response (Lewis et al., 2020). These innovations make it feasible for universities to implement reliable, context-aware digital assistants tailored to their regulatory environments.

In this context, the concept of an AI-assisted Regulations & Student Support Chatbot for the FoE at UoM emerges as a timely and strategic initiative. The proposed concept of the chatbot aims to bridge the gap between static regulatory documentation and students' real-time information needs by serving as an always-available virtual assistant capable of providing accurate and regulation-backed responses. Such a system would address questions on progression rules, module registration, assessment policies, late submission procedures, mitigating circumstances, credit limits, industrial training requirements, and other academic or administrative processes. By centralising information retrieval and enabling natural-language interaction, the chatbot would allow students to find relevant regulations without navigating lengthy documents or relying solely on staff availability.

Importantly, the objective of the chatbot is not to replace academic or administrative staff, but to optimise their workload by reducing repetitive queries and enabling them to focus on complex, high-impact student support. Similar implementations internationally have shown that AI support systems improve staff productivity, reduce administrative delays and enhance students' confidence in navigating university requirements (Okonkwo & Ade-Ibijola, 2021; Carayannopoulos, 2018). In resource-constrained contexts such as Mauritius, where institutions must balance increasing demands with limited human and technological infrastructure, AI-assisted systems offer particularly meaningful benefits.

Furthermore, this initiative aligns not only with broader educational and national strategies but also with the UoM strategic direction, namely digital transformation. Globally, institutions are encouraged to adopt AI responsibly to improve efficiency, transparency, and learner-centred services (Unesco, 2021). This specific UoM strategic direction emphasises digital transformation, student success and enhanced service delivery, making the proposed chatbot consistent with institutional priorities. Engineering students, who are often early adopters of emerging technologies, are an appropriate cohort to pilot such an innovation within the Mauritian higher education landscape.

The conceptual design of the chatbot incorporates principles of responsible AI usage, including data privacy, accuracy validation, and transparent sourcing of information. By employing a RAG-based architecture, the system can ensure that responses are grounded in official faculty and university regulations rather than generated solely from model inference. This mitigates the risk of hallucinations—incorrect yet confident AI-generated answers—which remains a concern in LLM-based systems (Ji et al., 2023). A governance framework for updating documents, maintaining version control, and ensuring quality assurance is also essential to ensure long-term system reliability.

In addition to providing regulatory information, the chatbot has the potential to support broader aspects of student success. Research shows that AI chatbots can act as low-barrier points of contact, particularly for students who may be hesitant to seek help from staff members (Følstad & Brandtzæg, 2017). They provide reminders, personalised guidance, and contextual explanations. This can positively influence student retention and engagement. Although the current concept focuses on programme regulations and administrative support, future extensions will include academic advising, learning resources, mental health signposting, or integration with learning management systems (LMS). Such scalability underscores the potential of the chatbot to evolve into a comprehensive digital companion throughout the student lifecycle.

1.1 Research aim

The aim of this study is to design and propose a governance-aware retrieval-augmented chatbot framework for supporting student access to academic regulations within a university environment. The system is intended to function as a first-line informational support tool, enabling students to query institutional policies while ensuring that responses remain grounded in authoritative regulatory documents and aligned with institutional governance requirements.

More specifically, the study aims to develop a conceptual system architecture and evaluation framework that integrates:

- a regulatory knowledge base constructed from official institutional documents,
- semantic information retrieval for locating relevant regulatory clauses,
- citation-based response generation to ensure traceability of information sources, and
- a risk-aware escalation mechanism that redirects high-risk queries to human administrators.

The overall objective is to propose how conversational AI systems can be designed to support transparent and accountable access to institutional regulations while respecting the governance boundaries of higher education institutions.

1.2 Research gap

Recent advances in LLMs and conversational AI have led to the rapid adoption of chatbot systems in higher education

for student support services. These systems are often used to answer questions related to admissions, course information, and administrative procedures (Kim & Baylor, 2016). However, existing implementations typically focus on general informational assistance and often rely on generative models that may produce responses without explicit grounding in institutional policy documents.

This raises several challenges when applying chatbot systems to regulatory guidance, where incorrect or unsupported responses may have academic, administrative, or legal consequences for students and institutions (Sharda et al., 2021).

Despite the growing use of conversational AI in education, several gaps remain in the literature:

1. Limited focus on regulatory governance

Most existing chatbot systems in higher education do not explicitly address how institutional regulations should be represented, retrieved, and cited within AI-generated responses.

2. Lack of risk-aware interaction models

Current chatbot implementations rarely distinguish between routine informational queries and high-risk cases such as appeals, disciplinary matters, or exceptional academic circumstances.

3. Insufficient attention to traceability and accountability

Many conversational AI systems prioritise fluent dialogue generation but provide limited mechanisms for tracing responses back to authoritative policy sources.

4. Absence of structured evaluation frameworks for regulatory AI systems

While chatbot performance metrics are widely studied in NLP, fewer studies propose evaluation frameworks specifically tailored to policy-grounded student advisory systems.

These gaps highlight the need for system designs that integrate information retrieval, governance constraints, and responsible AI practices when deploying conversational interfaces for institutional regulations (Holmes et al., 2019).

This paper introduces the conceptual framework for an AI-assisted Regulations & Student Support Chatbot for the FoE. It discusses the motivations, system architecture, expected benefits, ethical considerations, and implementation challenges. By conceptualising a tailored, context-aware and regulation-grounded AI system, this work can contribute to emerging scholarly discussions on AI-driven student support in the context of higher education in developing countries and offers a practical roadmap for enhancing student experience through responsible technological innovation.

2. Literature Review

The integration of AI in educational settings has garnered significant interest, particularly in enhancing student support systems and regulatory compliance. This literature review

examines the existing research on AI-assisted chatbots, their applications in educational institutions, and the implications for student engagement and academic success.

2.1 AI and educational technology

The advent of AI has transformed educational technology, enabling personalised learning experiences that cater to individual student needs. Luckin et al. (2016) emphasise the potential of AI to enhance learning by providing adaptive feedback and fostering critical thinking skills. Moreover, AI technologies facilitate the development of intelligent tutoring systems, which can offer real-time support and resources to students.

2.2 Chatbots as student support tools

AI chatbots have emerged as a viable solution for student support across various academic disciplines. Recent studies indicate that chatbots can effectively address common student inquiries, thus reducing the workload on administrative staff (De La Roca et al., 2024). For instance, the use of chatbots in universities has been shown to streamline processes such as course registration, scheduling, and answering frequently asked questions. A notable case study at Georgia State University demonstrated a significant increase in student retention rates after the implementation of a chatbot system (Nurshatayeva et al., 2021).

2.3 Regulatory compliance in education

Regulatory compliance is an essential aspect of higher education, ensuring that institutions meet established academic standards and practices. AI-driven solutions can assist in monitoring compliance by automatically tracking changes in regulations and facilitating timely updates to institutional policies. Duarte & Vardasca (2023) highlight the importance of integrating compliance systems within educational technology to maintain institutional integrity and support accreditation processes.

2.4 User experience and interaction design

The effectiveness of a chatbot is greatly influenced by its design, including user experience and interaction capabilities. A user-centric approach that prioritises intuitive interfaces and NLP enhances user engagement with chatbots. Furthermore, studies indicate that incorporating empathy and emotional intelligence within chatbot interactions can foster positive user experiences, leading to higher satisfaction among students (Jiang et al., 2022). The design of the chatbot for the FoE will prioritise these aspects to ensure effective communication and support.

2.5 Cultural considerations in chatbot development

When developing AI chatbots for educational contexts, it is crucial to consider cultural and contextual factors. Luo & Hsiao-Chin (2023) suggest that chatbots must be tailored to align with the linguistic and cultural nuances of their user base to enhance effectiveness. In the context of the UoM, understanding the local culture and communication styles will be pivotal in designing the chatbot that will resonate with students and staff alike.

2.6 Challenges and limitations

Despite the potential benefits of AI chatbots in education, challenges exist that must be addressed for successful implementation. Common barriers include resistance to change from staff and students, concerns over data privacy, and the need for ongoing maintenance and updates (Yin, 2021). It is essential for the UoM to develop a comprehensive strategy that includes training for users and mechanisms for continuous improvement of the chatbot.

2.7 Future directions

Looking forward, the evolution of AI in education suggests a shift towards more integrated systems that combine chatbots with other technological innovations, such as LMS and data analytics tools. By enabling seamless interactions among various platforms, higher educational institutions can create a comprehensive support ecosystem that enhances both academic and administrative functions (Ashok et al., 2021).

The literature underscores the transformative potential of AI-assisted chatbots in enhancing student support and regulatory compliance within higher educational institutions. By leveraging insights from existing research, the proposed chatbot for the FoE aims to foster a supportive academic environment that aligns with current technological advancements and meets the needs of its diverse student body. Future research should continue to explore the impacts of such technologies and iterate on their designs to maximise effectiveness and user engagement.

3. System Model

3.1 Study type

This research adopts a design-oriented conceptual study that proposes a system architecture and evaluation framework for a regulatory-support chatbot.

The study does not present a deployed system or empirical performance evaluation. Instead, it focuses on:

- defining the structure of a regulatory knowledge base,
- proposing a retrieval-augmented chatbot architecture, and
- outlining a set of evaluation metrics and governance mechanisms for assessing such systems.

The work can therefore be characterised as a design science research contribution, where the primary output is a conceptual artefact consisting of:

- a system architecture,
- a dataset design methodology, and
- an evaluation framework for regulatory chatbot systems.

By articulating the design principles, implementation considerations and evaluation criteria, the study aims to provide a foundation for future empirical implementations and experimental evaluations of AI-assisted regulatory support tools in higher education.

3.2 Expected contribution

The contribution of this study lies in the conceptualisation of a governance-aware chatbot framework tailored to institutional regulations in higher education. Specifically, the study contributes:

1. A structured approach to constructing a regulatory knowledge base suitable for conversational AI systems.
2. A retrieval-augmented system architecture designed to improve traceability and citation of regulatory sources.
3. A risk-aware query classification and escalation model to ensure appropriate handling of sensitive student queries.
4. A proposed evaluation framework for assessing the reliability, governance compliance, and usability of regulatory-support chatbots.

Together, these elements provide a methodological foundation that can support future research and practical deployments of AI systems for regulatory guidance in higher education institutions.

3.3 System overview

The proposed AI-assisted Regulations and Student Support Chatbot is designed to provide reliable, regulation-grounded guidance to students of the FoE. The system will prioritise accuracy, traceability, and accountability by grounding all responses in official institutional documents and embedding governance mechanisms that enable human oversight.

To achieve these objectives, the system adopts a Retrieval-Augmented Conversational Architecture, where semantic retrieval from a curated regulatory knowledge base precedes controlled response generation.

3.4 Block diagram of the proposed system

The block diagram of the proposed system is shown in Figure 1.

3.4.1 User interaction module

This module will provide the interface through which students interact with the chatbot using natural language queries. The interface will be accessible via web and mobile platforms and supports bilingual interaction in English and French. The module will ensure ease of use, accessibility, and clear communication of the system's advisory role.

3.4.2 Query processing and risk classification module

Once a query is received, it will be processed to identify its language, remove irrelevant noise, and standardise the input. The query will then be analysed to determine its risk category, distinguishing between routine informational requests (e.g., credit requirements, deadlines) and high-stakes regulatory matters (e.g., appeals, disciplinary issues).

This risk-aware classification plays a critical role in downstream decision-making, particularly in determining confidence thresholds and escalation requirements. A flowchart to visualise the user's journey from query to response is shown in Figure 2.




Figure 1: Block diagram of the AI-assisted Regulations and Student Support Chatbot.




Figure 2: Flow chart on User's journey from query to response.

3.4.3 Knowledge base and document management module

This module will represent the authoritative backbone of the system. It will contain official documents issued by the FoE, including:

- general academic regulations,
- faculty and programme handbooks,
- programme regulations,
- assessment and progression rules, and
- official circulars and notices.

The relevant documents will be segmented into structured regulatory units (clauses or subsections) and enriched with metadata such as document source, section number, and validity period. A version control mechanism ensures that outdated regulations are archived and that only current policies are used for active responses.

3.4.4 Information retrieval module

The Information Retrieval Module will perform semantic matching between the processed user query and the regulatory units stored in the knowledge base. Instead of relying solely on keyword matching. This module will identify conceptually relevant clauses, enabling effective handling of informal or paraphrased student queries.

Only the most relevant regulatory segments will be forwarded to the response generation stage, ensuring precision and relevance.

3.4.5 Retrieval-augmented response generation module

This module will generate student-facing responses using an RAG approach. The retrieved regulatory segments will be used as authoritative context, constraining the response to verified information.

The generated response aims to:

- answer the query clearly and concisely,
- explain applicable regulations in accessible language, and
- reference the originating regulatory document and section.

If sufficient regulatory information is unavailable, the system will explicitly state its limitation rather than generating speculative advice.

3.4.6 Response validation and confidence estimation module

Before delivering a response, the system will evaluate its reliability using a confidence estimation mechanism. This evaluation considers:

- relevance and consistency of retrieved documents,
- completeness of the regulatory context, and
- the identified risk level of the query.

A confidence score will be computed and compared against a predefined threshold to determine whether the response can be safely delivered or requires escalation.

3.4.7 Human escalation and governance module

For queries classified as high risk or responses with insufficient confidence, the system will activate a human escalation pathway. In such cases, the chatbot will provide general guidance and direct the student to the appropriate Faculty authority (e.g., Programme Coordinator, Faculty Registry, Examinations Office).

This module will also support governance functions such as interaction logging, auditability, and system performance monitoring, ensuring compliance with institutional policies and ethical guidelines.

3.5 Integration of the system model

The interaction between the functional blocks will ensure that:

- all responses are regulation-grounded,
- high-risk decisions remain human-controlled,
- transparency and traceability are maintained, and
- the system remains adaptable to regulatory changes.

By embedding governance and escalation mechanisms directly into the system architecture, the proposed model will be designed to align with best practices for deploying AI systems in regulated educational environments.

3.6 Summary

The proposed system architecture will adopt a risk classifier coupled with a human-in-the-loop (HITL) decision layer, grounded in established principles from decision support systems (DSS) and AI ethics. In DSS literature, AI systems are explicitly designed not as autonomous decision-makers but as tools that support human judgment, providing recommendations that must be validated or overridden by a human expert. For example, Power (2002) emphasises that DSS are intended to augment, rather than replace, human decision-making, particularly in semi-structured and unstructured problem domains.

The inclusion of a risk classifier is motivated by the need to operationalise uncertainty and decision criticality. By estimating the likelihood that a response may be incorrect, ambiguous, or context-dependent, the system can distinguish between low-risk queries (e.g., factual FAQs) and high-risk cases (e.g., policy interpretation or academic decisions). This aligns with research in human-AI systems design, where Horvitz (1999) introduced principles of mixed-initiative interaction, advocating for systems that dynamically determine when to act autonomously and when to defer to human control. Hypothetically, effective deployment would result in high automation accuracy for routine queries while selectively escalating complex cases, thereby balancing efficiency and reliability.

From an AI ethics perspective, the HITL approach is essential to ensure accountability, fairness, and oversight. The European Commission's High-Level Expert Group on Artificial Intelligence (European Commission, 2019) explicitly identifies human agency and oversight as a core requirement for trustworthy AI systems. Similarly, Dignum (2019) argues that responsible AI must incorporate mechanisms for human control, particularly in high-stakes or socially sensitive contexts such as education. In practice,

this means that the system should not only generate answers but also provide confidence estimates and trigger escalation mechanisms when thresholds are not met. A desirable outcome would be a system where human oversight is concentrated on genuinely complex or high-stakes cases, while routine interactions are safely automated.

Finally, the integration of explainability within this architecture reflects broader regulatory and ethical trends. Explainable AI is increasingly viewed as a prerequisite for trustworthy DSS, particularly in domains where users must justify or rely on AI-supported decisions. Doshi-Velez & Kim, (2017) highlight the importance of interpretability for ensuring transparency, user trust, and effective human-AI collaboration. In this context, the combination of retrieval grounding, risk classification, and human oversight forms a coherent socio-technical design: the system provides evidence-based responses, evaluates its own uncertainty, and defers to human judgment when necessary. This layered approach is therefore well-justified both theoretically and practically for an AI system intended to support student success and institutional decision-making.

The block diagram and accompanying system model illustrate a modular, governance-aware AI chatbot architecture tailored to the regulatory and operational context of the FoE. The design balances automation with institutional oversight, ensuring that the system enhances student support while preserving academic integrity and regulatory compliance.

4. Discussion

4.1 Overview of the system model in practice

The proposed system model proposes how an AI-assisted chatbot can be safely deployed in a highly regulated academic environment such as the FoE. Unlike general-purpose conversational agents, this system will intentionally be constrained by design choices that prioritise regulatory correctness, traceability, and institutional accountability.

The layered architecture—comprising query processing, semantic retrieval, retrieval-augmented response generation, validation, and human escalation—ensures that automation is applied selectively. This design directly addresses concerns identified in the literature regarding hallucinations, liability, and loss of trust in AI-driven student services. The discussion that follows critically examines how each component of the system model influences dataset design and metric selection as reflected in the evaluation tables.

4.2 System implementation details

This section describes the implementation of the regulatory-support chatbot system, including the architecture, major design decisions, conflict-resolution mechanisms, and potential failure modes. The objective of this description is to enable reproducibility of the system design and facilitate replication in similar institutional environments.

4.2.1 System architecture

The system follows a RAG architecture composed of five primary modules:

1. Document ingestion and preprocessing module,
2. Vector indexing and metadata storage,
3. Query processing and semantic retrieval module,
4. Risk-aware classification and escalation module, and
5. Response generation and citation module.

Each component is described below.

4.2.1.1 Document ingestion and preprocessing

Document Sources

The system ingests official institutional documents from faculty and university repositories, including:

- university regulations,
- faculty regulations,
- programme handbooks,
- assessment rules, and
- circulars and notices.

Documents are processed only if they originate from verified institutional sources.

Text extraction

Documents in PDF or DOCX format are converted into plain text using automated parsing tools. Structural markers such as headings, section numbers, and clause identifiers are preserved to maintain the legal context of regulations.

Clause-level segmentation

Documents are segmented into regulatory clauses, which serve as the fundamental retrieval unit.

Segmentation follows these rules:

1. Each numbered rule is treated as a separate clause.
2. Bullet-point policy statements are treated as independent clauses.
3. Long paragraphs are split into sentences when they contain multiple regulatory conditions.

Each clause is stored as a structured record containing:

- clause text,
- document identifier,
- section number,
- document category,
- language,
- faculty applicability, and
- validity period.

This segmentation decision was made to improve semantic retrieval precision. Clause-level indexing reduces ambiguity when answering specific regulatory queries.

4.2.1.2 Vector indexing and knowledge base construction

Each clause is embedded into a semantic vector representation using a pre-trained sentence embedding model.

The embedding process converts each clause into a high-dimensional vector that captures semantic meaning. These vectors are stored in a vector index that supports approximate nearest-neighbour search.

Index structure

The knowledge base contains two linked storage layers:

Vector index

- stores clause embeddings, and
- used for semantic similarity search.

Metadata database

- stores document metadata and governance fields, and
- used for filtering and conflict resolution.

The metadata database is linked to the vector index through a unique clause identifier.

Metadata filtering

During retrieval, candidate clauses are filtered based on:

- faculty applicability,
- programme applicability, and
- validity period.

This ensures that obsolete or irrelevant regulations are not returned in responses.

4.2.1.3 Query processing pipeline

When a user submits a query, the system performs the following sequence of operations.

Step 1: Query normalisation

The input query is cleaned and standardised through:

- removal of formatting artefacts,
- language detection, and
- token normalisation.

If the query is written in French, a multilingual embedding model is used to ensure semantic compatibility with English documents.

Step 2: Risk classification

The query is classified into one of three risk levels:

- low risk,
- medium risk, and
- high risk.

Risk classification is performed using a rule-based classifier supported by keyword patterns and query intent analysis.

Examples of high-risk indicators include terms such as:

- appeal,
- disciplinary,
- misconduct,
- mitigation, and
- exemption request.

High-risk queries are flagged for potential escalation.

Step 3: Semantic retrieval

The query embedding is compared with clause embeddings stored in the vector index.

The system retrieves the **top-k most similar clauses**, typically between three and five candidates.

Similarity is computed using cosine similarity between vectors.

Candidate clauses are then filtered using metadata constraints.

Step 4: Evidence ranking

Retrieved clauses are ranked according to:

- semantic similarity score,
- document authority level, and
- regulatory recency.

Authority level is determined using the following hierarchy:

1. university regulations,
2. faculty regulations,
3. programme handbooks, and
4. circulars and notices.

Higher-authority documents are prioritised when multiple clauses match a query.

4.2.1.5 Response generation and citation

The final response is generated by a language model conditioned on the retrieved clauses.

The prompt structure includes:

- the original query,
- the retrieved clauses,
- metadata about each clause, and
- instructions to cite the source document.

The model is instructed to:

1. summarise the relevant regulations,
2. provide citations referencing document title and section number, and
3. avoid making unsupported claims outside the retrieved evidence.

If insufficient evidence is retrieved, the system produces a fallback message requesting clarification or recommending contact with administrative staff.

4.2.1.5 Handling conflicting regulations

Conflicts between regulatory documents may occur when:

- older regulations remain accessible,
- programme-specific rules differ from general faculty rules, and
- circulars temporarily modify official policies.

To address this issue, the system implements a rule precedence hierarchy.

Regulation precedence order

When conflicting clauses are retrieved, the following precedence rules are applied:

1. Latest valid regulation

Clauses with a more recent validity period override older clauses.

2. Higher authority level

University-wide regulations override faculty or programme rules.

3. Programme-specific rules

Programme rules override generic faculty rules when applicable.

4. Temporary circulars

Circulars override regulations only if explicitly marked as active.

Conflict detection

A conflict is detected when retrieved clauses contain contradictory statements, such as different credit requirements or deadlines.

When conflicts are detected, the system performs one of the following actions:

- present both clauses with clarification,
- prioritise the clause with higher precedence, and
- escalate the query to human administration.

The escalation option is used when the system cannot determine a clear resolution.

4.3 Failure modes

Several failure modes were identified during system design.

Retrieval failure

Relevant clauses may not be retrieved due to:

- incomplete document coverage,
- semantic mismatch between query and clause wording, and
- insufficient embedding representation.

Mitigation strategies include expanding the document corpus and tuning the retrieval model.

Hallucinated responses

The language model may generate unsupported statements not present in the retrieved evidence.

To reduce this risk, the system restricts responses to retrieved clauses and requires explicit citation.

Metadata errors

Incorrect metadata annotation (e.g., wrong validity period) may cause outdated regulations to appear in responses.

Periodic dataset validation is required to maintain metadata accuracy.

Risk misclassification

A query may be incorrectly classified as low-risk when it requires administrative oversight.

To mitigate this, the system adopts a conservative strategy where ambiguous queries are escalated.

4.4 Replication requirements

To replicate the implementation described in this study, the following components are required:

- a corpus of institutional regulatory documents,
- a clause-level segmentation pipeline,
- a multilingual embedding model,
- a vector search index supporting semantic retrieval,
- a metadata database containing governance attributes,
- a LLM for response generation, and
- a rule-based risk classification system.

Together, these components form a reproducible framework for building regulatory-aware chatbot systems in higher education environments.

4.5 Alignment between system architecture and dataset design

A central strength of the system model lies in the tight coupling between architecture and dataset construction. The regulatory document corpus, described in [Table 1](#) (dataset composition summary), will not be treated as unstructured text but as a curated, version-controlled institutional knowledge base.

[Table 1](#) categorises the dataset into document types (e.g., general regulations, faculty handbooks, programme specifications, assessment regulations) and reports attributes such as:

- number of documents per category,
- average number of regulatory clauses per document,
- language distribution (English/French), and
- update frequency.

[Table 1](#) highlights that document granularity directly affects retrieval performance. Clause-level segmentation enables the Information Retrieval Module to achieve higher Recall scores compared to paragraph-level segmentation, as it reduces ambiguity and improves semantic matching. However, finer segmentation also increases index size and retrieval complexity, which must be managed through efficient indexing strategies.

Table 1: Dataset Composition Summary (Regulatory Knowledge Base)

Document Category	Language	Update Frequency
General University Regulations	English	Annual
Faculty of Engineering Regulations	English	Annual
Programme Handbooks (BEng/MSc)	English	Annual
Assessment & Examination Rules	English	Biennial
Faculty Circulars & Notices	English / French	Ad hoc

The inclusion of metadata fields such as validity period and faculty applicability—outlined in Table 1—supports the system’s governance objectives and prevents outdated or cross-faculty regulations from being surfaced to students.

4.5.1 Regulatory knowledge base dataset construction and metrics

The regulatory knowledge base dataset was constructed to represent the institutional rules and policies governing academic administration within the FoE. The objective of the dataset design is to provide a structured and traceable source of regulatory information that could support reliable retrieval and citation by the chatbot system.

Source documents:

The dataset was compiled from official institutional documents obtained from university repositories and faculty administrative records. The sources include:

- general university regulations and academic rules,
- FoE regulatory documents,
- programme handbooks for undergraduate and postgraduate programmes,
- assessment and examination regulations, and
- faculty circulars and administrative notices.

Only authoritative documents formally issued by the university or faculty were included in the dataset to ensure that the system references official policy sources. Informal communications such as emails or student forum posts were excluded.

Document processing pipeline:

The dataset construction process followed a multi-stage pipeline:

1. Document collection:

Regulatory documents were collected in digital format (primarily PDF and Word documents). Each document was assigned a unique document identifier.

2. Text extraction:

Text content was extracted from the documents using automated parsing tools. Non-textual elements such as images and decorative formatting were removed. Section headings, numbering structures, and clause identifiers were preserved to maintain regulatory context.

3. Clause-level segmentation:

To improve retrieval precision, documents were segmented into regulatory clauses, defined as the smallest

self-contained unit of policy that expresses a rule or condition.

A clause was operationally defined as:

- a numbered rule,
- a bullet-point policy statement, and
- or a sentence expressing a specific requirement or condition.

Each clause was stored as an independent retrieval unit while maintaining links to its parent document and section.

Clause segmentation was selected instead of paragraph-level segmentation because it reduces semantic ambiguity and enables more precise retrieval of policy statements.

4. Metadata annotation:

Each clause entry was enriched with metadata fields to support governance-aware retrieval and filtering. The metadata schema included:

Metadata Field	Description
Document Category	Type of institutional document (e.g. regulations, handbook, circular)
Faculty Applicability	Faculty or institutional scope
Programme Applicability	Programme or degree to which the clause applies
Language	Language of the clause (English or French)
Validity Period	Time interval during which the regulation is applicable
Update Frequency	Expected revision cycle of the document
Source Reference	Original document title and section number

The inclusion of validity period and faculty applicability ensures that outdated or irrelevant policies are not retrieved during query processing.

Dataset attributes

The composition of the regulatory dataset uses the following metrics:

Number of documents per category:

The total count of documents belonging to each document category.

Formally:

Number of Documents (category i) = count of all documents assigned to category i

This metric provides an overview of dataset coverage across different regulatory sources.

Average number of regulatory clauses per document

The number of clause-level entries extracted from documents within each category.

Formula:

Average clauses per document = (Total number of extracted clauses in category) ÷ (Number of documents in category)

This metric reflects document complexity and influences the number of indexed retrieval units.

Large distribution:

The proportion of clauses written in each supported language.

Language Proportion (L) = (Number of clauses in language L) ÷ (Total number of clauses)

This metric helps determine the extent of multilingual content in the regulatory dataset.

Update frequency

The expected revision cycle of documents within a category. This value is derived from institutional policy review schedules.

Typical categories include:

- annual updates (e.g., programme handbooks),
- biennial updates (e.g., examination regulations), and
- ad hoc updates (e.g., circulars and notices).

Update frequency informs dataset maintenance strategies and helps determine when re-indexing of the knowledge base is required.

Impact of document granularity

Document granularity refers to the size of the textual units used for retrieval. In this dataset, clause-level segmentation is used as the primary indexing unit.

Smaller retrieval units reduce semantic ambiguity by isolating specific rules or requirements. However, finer segmentation increases the total number of indexed items, which may increase storage requirements and retrieval computation. Efficient indexing structures (such as vector embeddings and metadata filtering) are therefore required to maintain retrieval efficiency.

4.5.2 Query dataset construction and risk classification

A separate query dataset will be created to simulate realistic student interactions with the chatbot system and to evaluate system behaviour under different levels of institutional risk.

Query source generation

Queries will be generated using three complementary approaches:

1. Student enquiry logs

Common administrative questions collected from faculty administrative offices.

2. Programme handbook scenarios

Hypothetical questions derived from rules and conditions described in programme regulations.

3. Expert-generated queries

Queries formulated by academic administrators to represent complex regulatory cases.

4. The resulting query set will be designed to cover the most common regulatory topics encountered by students.

Query categories

Queries will be grouped into thematic categories reflecting major administrative processes:

- programme progression,
- assessment and examinations,
- registration and enrolment,
- graduation requirements,
- appeals and mitigation, and
- disciplinary or special cases.

These categories correspond to major regulatory domains within faculty administration.

Risk level classification

Each query will be assigned a risk level representing the potential consequences of providing an incorrect automated response. Risk classification will be performed using the following criteria:

Low-risk queries

Queries requesting informational clarification about regulations.

Examples include:

- module registration procedures,
- examination schedules, and
- general progression rules.

Errors in these responses are unlikely to create institutional or legal consequences.

Medium-risk queries

Queries related to eligibility conditions or academic outcomes.

Examples include:

- graduation eligibility,
- credit completion requirements, and
- programme progression thresholds.

Incorrect responses could mislead students but typically remain correctable through administrative follow-up.

High-risk queries

Queries involving appeals, disciplinary actions, or exceptional cases.

Examples include:

- academic appeals,
- misconduct procedures, and
- requests for special mitigation.

These cases often require formal institutional review and human oversight, and therefore must trigger escalation rather than automated decision-making.

4.6 Query dataset characteristics and risk distribution

Table 2 (query dataset characteristics) summarises the evaluation query set used to test the system. This table typically includes:

- total number of queries,
- thematic categories (progression, assessment, appeals, graduation, enrolment),
- proportion of high-risk versus low-risk queries, and
- language distribution.

4.7 Interpretation of retrieval performance metrics

Retrieval performance metrics presented in Table 3 (information retrieval performance) would be useful in evaluating how effectively the system identifies relevant regulatory clauses. The table is typically expected to report:

- Top-1, Top-3, and Top-5 accuracy,
- Recall values across query categories.

Table 3 is designed to help assess the semantic retrieval module performance. The findings can thus help inform future regulatory drafting and consolidation efforts at the institutional level.

4.8 Discussion of response accuracy and grounding metrics

Table 4 (response quality and grounding metrics) presents expert-evaluated measures such as:

Table 2: Query Dataset Characteristics

Query Category	Risk Level
Programme Progression	Low
Assessment & Exams	Low
Registration & Enrolment	Low
Graduation Requirements	Medium
Appeals & Mitigation	High
Disciplinary / Special Cases	High

Table 3: Information Retrieval Performance

Metric	Top-1	Top-3	Top-5
Retrieval Accuracy (%)	–	–	–
Recall (%)	–	–	–

Table 4: Response Quality and Grounding Metrics (Expert Evaluation)

Metric	Value (%)
Regulatory Accuracy Rate	–
Citation Correctness	–
Hallucination Rate	–
Partial / Ambiguous Responses	–

- regulatory accuracy rate,
- citation correctness, and
- hallucination rate.

Hallucination rate measures the extent to which the system generates information that is not supported by any source document. A low hallucination rate would indicate that the RAG approach is effectively grounding responses in verified content, whereas a high rate would suggest that the model is still relying too heavily on its generative component, increasing the risk of misinformation.

Citation correctness evaluates whether the references provided by the system accurately point to the relevant sections of the source documents. High citation correctness would imply that the system not only retrieves appropriate information but also links it transparently to its origin, thereby enhancing user trust. Conversely, low citation correctness would indicate mismatches between answers and their cited sources, which could undermine confidence even if the answer itself appears plausible.

Answer accuracy assesses the factual correctness and completeness of the generated response relative to the query. Strong performance on this metric would suggest that the system can correctly interpret and synthesise information from retrieved documents, while weaker performance would highlight limitations in reasoning or comprehension, particularly for complex or multi-part queries.

Grounding completeness measures how fully the response is supported by the retrieved evidence, especially in cases where multiple clauses or documents must be combined. High grounding completeness would indicate that the system successfully integrates all relevant pieces of information, whereas partial scores would suggest that some necessary context is missing, potentially leading to incomplete or oversimplified answers.

Finally, escalation or confidence calibration metrics capture the system’s ability to recognise its own uncertainty and defer to human intervention when necessary. Effective calibration would be reflected in the system flagging ambiguous or low-confidence cases for review, thereby reducing the risk of misleading outputs. Poor calibration, on the other hand, would result in overconfident responses in situations that require nuanced interpretation, particularly when regulations span multiple clauses or contain inherent ambiguities.

4.9 Risk handling and escalation performance

Table 5 (risk handling and escalation metrics) evaluates the system’s governance behaviour through:

Table 5: Risk Handling and Escalation Metrics

Metric	Value (%)
Correct Escalation Rate	–
False Automation Rate	–
Unnecessary Escalation Rate	–

- correct escalation rate,
- false automation rate, and
- unnecessary escalation rate.

Correct escalation rate measures the proportion of cases where the system appropriately identifies that it cannot confidently or safely provide a fully reliable answer and escalates the query to a human. A high correct escalation rate would suggest that the system is effectively recognising its limitations, particularly in complex, ambiguous, or high-stakes scenarios. A low rate, on the other hand, would indicate that the system is failing to flag cases that genuinely require human judgment, increasing the risk of incorrect or misleading responses being delivered automatically.

False automation rate captures the proportion of cases where the system chooses to provide an automated response when it should have escalated instead. This is a critical risk metric, as high false automation implies overconfidence: the system is answering queries autonomously despite insufficient evidence, ambiguity, or low certainty. A low false automation rate would indicate safer behaviour, where the system avoids making unsupported decisions and defers appropriately when needed.

Unnecessary escalation rate measures how often the system escalates queries that it could have handled correctly on its own. A high unnecessary escalation rate would suggest that the system is overly conservative, potentially leading to inefficiencies, increased workload for staff, and reduced user satisfaction due to delays. Conversely, a low unnecessary escalation rate would indicate that the system is making efficient use of automation, handling straightforward queries independently while reserving escalation for genuinely difficult or uncertain cases.

Together, these three metrics provide a balanced view of system behaviour: correct escalation reflects appropriate caution, false automation highlights risky overconfidence, and unnecessary escalation captures inefficiency due to excessive caution. An effective system should aim for high correct escalation, low false automation, and low unnecessary escalation.

4.10 User-centric evaluation and perceived trust

Table 6 (user experience and trust metrics) gives visibility on survey-based measures including:

- User satisfaction score,
- Perceived trustworthiness, and
- Query resolution rate.

The discussion of **Table 6** will help reveal how trustworthiness is correlated with citation transparency and response verbosity

Table 6: User Experience and Trust Metrics (Student Survey)

Metric	Mean Score
User Satisfaction (1–5)	–
Perceived Trustworthiness (1–5)	–
Clarity of Responses (1–5)	–
Resolution Rate (%)	–
Willingness to Reuse (%)	–

Scale: 1 = Very Poor, 5 = Excellent

or conversational fluency. This finding will validate the system model's prioritisation of traceability over purely natural dialogue.

4.11 System model limitations revealed by evaluation

The discussion will also critically examine limitations exposed by the metrics and datasets:

1. **Regulatory ambiguity:** Some regulations are inherently open to interpretation, limiting full automation.
2. **Dataset bias:** The query dataset reflects common enquiries and may under-represent rare but critical edge cases.
3. **Temporal dependency:** Evaluation results depend on the currency of regulations; performance may degrade if updates are delayed.

These limitations are architectural rather than algorithmic and reinforce the need for continuous governance and institutional oversight.

4.12 Implications for the FoE and broader adoption

From an institutional perspective, the discussion aims to illustrate that the proposed system model is not merely a technical artefact but an operational support mechanism. The evaluation tables collectively will provide indications on how the chatbot will:

- possibly reduce routine administrative workload,
- possibly improve consistency of regulatory advice, and
- possibly enhance student confidence in institutional processes.

The modular design will also enable future extensions, such as integration with LMS or analytics-driven policy refinement. The effectiveness is a theoretical expectation until the pilot phase is complete.

4.13 Ethical considerations, consent procedures, and data management plan

4.13.1 Ethical approval governance

The development and evaluation of the regulatory-support chatbot will be designed to comply with institutional research ethics guidelines and applicable data protection regulations. Because the system processes student queries and institutional

regulatory documents, the study incorporates safeguards to ensure that personal data, institutional information, and user interactions are handled responsibly.

Prior to dataset construction and system evaluation, ethical approval would be sought from the relevant institutional research ethics committee. The ethical review would assess the handling of student interaction data, the anonymisation procedures used during dataset preparation, and the mechanisms implemented to prevent the disclosure of personally identifiable information.

The study is designed to minimise ethical risks by ensuring that:

- regulatory documents originate from publicly available or institutionally authorised sources,
- student queries used in the evaluation dataset are anonymised and stripped of personal identifiers, and
- the chatbot system does not provide binding administrative decisions and instead functions as an informational support tool.

4.13.2 Consent procedures

4.13.2.1 Student interaction data

Given that real student queries will be used to construct the evaluation dataset, informed consent procedures will be implemented.

Students participating in the study would be informed that:

- their interactions may be recorded for research and system improvement purposes,
- all collected data will be anonymised before analysis, and
- participation is voluntary and may be withdrawn at any time.

Consent would be obtained through a digital consent form prior to participation. The consent form would clearly describe:

- the purpose of the research,
- the types of data collected,
- how the data will be stored and processed,
- potential risks and benefits of participation, and
- contact details of the research team and ethics committee.

Students who do not provide consent would not have their queries included in the research dataset.

4.13.2.2 Administrative staff participation

Where administrative staff contribute example queries or assist in validating regulatory interpretations, their participation is treated as expert input rather than personal data collection. Staff contributors would be informed that their input may be used for system design and evaluation purposes.

4.13.3 Data anonymisation

All user-generated queries included in the dataset undergo anonymisation before storage or analysis.

The anonymisation process removes or masks the following elements:

- student names,
- identification numbers,
- email addresses,
- programme registration numbers, and
- any other personally identifiable information.

Queries are stored as de-identified textual entries, and each query is assigned a random identifier to prevent linkage to individual students.

4.13.4 Data Management Plan (DMP)

A structured DMP was developed to ensure that the research data are securely stored, responsibly used, and reproducible where appropriate.

4.13.4.1 Data types

The project will produce and use several categories of data:

1. Regulatory document dataset

- university regulations
- faculty regulations
- programme handbooks
- examination rules
- circulars and notices

2. Clause-level regulatory dataset

- segmented policy clauses
- associated metadata

3. Query dataset

- anonymised student queries
- synthetic queries generated by experts

4. System evaluation logs

- retrieved clauses
- system responses
- escalation decisions

4.13.4.2 Data storage

All research data are stored in secure institutional storage systems with controlled access.

Security measures include:

- encrypted storage servers,
- password-protected access, and
- restricted access to authorised researchers only.

Raw datasets containing potentially sensitive information are stored separately from processed research datasets.

4.13.4.3 Data retention

Data will be retained for a limited period to support research validation and replication.

Typical retention policy:

- anonymised datasets: retained for 5 years,
- raw interaction logs: retained for 1–2 years, and
- temporary processing files: deleted after preprocessing.

Retention policies follow institutional data governance guidelines.

4.13.4.4 Data sharing and reproducibility

To support research transparency and reproducibility, selected components of the dataset may be shared publicly, subject to institutional approval.

The following resources may be released:

- anonymised query dataset,
- clause-level regulatory dataset structure,
- metadata schema, and
- system architecture description.

However, full regulatory documents may not be redistributed if they are subject to institutional copyright or restricted access policies.

Researchers wishing to replicate the system can reconstruct the knowledge base using publicly available institutional regulations.

4.13.4.5 Data security and privacy protection

The system is designed to minimise the risk of exposing sensitive information.

Key safeguards include:

- removal of personal identifiers from query data,
- role-based access control for dataset access, and
- separation of research datasets from operational system logs.

Additionally, the chatbot system itself is configured not to store long-term conversation histories unless explicitly required for evaluation purposes.

4.13.4.6 Data lifecycle

The research data lifecycle consists of the following stages:

1. data acquisition (regulatory documents and anonymised queries),
2. data preprocessing and clause segmentation,
3. metadata annotation and vector indexing,
4. system evaluation using the query dataset,
5. secure storage and documentation of datasets, and
6. controlled sharing of anonymised datasets for replication.

4.13.5 Limitations and ethical safeguards

Although the system will be designed to support regulatory information access, it is not intended to replace official administrative advice. All responses generated by the chatbot include a disclaimer indicating that:

- the response is informational,
- official confirmation should be obtained from faculty administration for binding decisions.

High-risk queries such as appeals, disciplinary matters, or exceptional cases are automatically escalated to human administrators to ensure appropriate governance oversight.

4.13.6 Compliance with data protection regulations

The study design aligns with principles commonly found in international data protection frameworks, including:

- data minimisation,
- purpose limitation,
- informed consent, and
- secure storage of personal data.

Only the minimum necessary data required for research evaluation will be collected and processed.

4.14 Summary of discussion

The discussion confirms that the proposed system model, when evaluated using carefully designed datasets and metrics, will potentially achieve a balanced integration of automation and governance. The tables on datasets and metrics are not only evaluative tools but also diagnostic instruments that reveal strengths, limitations, and opportunities for institutional improvement.

Overall, the system model is shown to be technically sound, institutionally aligned, and ethically grounded, making it suitable for deployment within the FoE and adaptable to similar higher education contexts.

5. Conclusion and Future Work

5.1 Conclusion

This paper presented the conceptual design, system model, and evaluation framework for an AI-assisted Regulations and Student Support Chatbot tailored to the FoE at the UoM. Unlike generic conversational systems, the proposed chatbot will explicitly be designed for high-stakes regulatory guidance, where accuracy, traceability, and institutional accountability are paramount.

The system model integrates semantic information retrieval, retrieval-augmented response generation, confidence-based validation, and human escalation mechanisms into a cohesive, governance-aware architecture. By grounding all responses in official University and Faculty regulations, the chatbot mitigates the risk of hallucinated or misleading information—an issue frequently cited in the literature on LLMs in education.

The implementation strategy emphasises modularity and sustainability, enabling the system to be maintained and updated in line with evolving academic regulations. The dataset construction process, based on authoritative regulatory documents and representative student queries, ensures contextual relevance and institutional alignment.

Tables 1 to 6 present the proposed evaluation framework designed to assess the performance and governance compliance of the system as a first-line student support tool. The framework includes metrics related to retrieval accuracy, regulatory correctness, hallucination rate, and citation correctness, which will be intended to evaluate how effectively the system identifies, interprets, and references relevant institutional regulations. These indicators are included to ensure that responses remain grounded in authoritative sources and aligned with institutional policies.

In addition, the framework incorporates measures for evaluating the behaviour of the risk-aware escalation strategy. Metrics such as correct escalation rate and false automation rate are defined to assess whether high-risk or policy-sensitive queries will be appropriately redirected to human administrators. These measures aim to support the monitoring of institutional boundaries and ensure that automated responses will be limited to suitable contexts.

User-centred evaluation metrics are also included to capture student perceptions of the system. Indicators such as satisfaction, perceived trustworthiness, and intention to reuse are proposed to assess how students experience and interact with the chatbot interface. The inclusion of these measures reflects the design emphasis on transparency, traceability, and the provision of clearly cited regulatory information rather than purely conversational fluency.

Overall, the evaluation framework is intended to provide a structured approach for assessing how a regulatory-support chatbot could operate within the FoE. The framework aims to support future empirical studies that examine regulatory clarity, administrative workload, and the student information experience while maintaining alignment with institutional governance and academic integrity requirements.

5.2 Limitations

Despite its strengths, the proposed system has several limitations. First, regulatory language is inherently complex and, in some cases, open to interpretation, which constrains full automation. Second, the evaluation dataset, while representative, may not capture rare or exceptional edge cases. Third, system performance is dependent on the timeliness and completeness of regulatory updates; delays in document integration could affect response accuracy.

These limitations are primarily institutional and procedural rather than algorithmic, underscoring the importance of sustained governance and human oversight.

5.3 Future work

Several avenues for future work emerge from this conceptual framework:

1. Functional pilot study

A controlled pilot will be conducted within selected engineering modules to deploy the chatbot and analytics platform in a live academic setting. This phase will focus on:

- testing system integration with LMS and institutional workflows,
- monitoring real-time student interactions with the chatbot,
- evaluating usability, response quality, and system reliability, and
- collecting feedback from both students and academic staff.

Hypothetically, successful outcomes would include high user engagement, reduction in routine queries handled by staff, and positive user satisfaction scores.

2. Empirical validation

Following the pilot, a rigorous empirical evaluation will be conducted to measure the system's impact on student success and decision support effectiveness. This will include:

- quantitative analysis of pass rates, grade improvements, and retention metrics,
- evaluation of predictive model performance (precision, recall, AUC),
- assessment of escalation metrics (correct escalation, false automation, unnecessary escalation), and
- comparative analysis using control vs intervention groups (A/B testing).

Strong empirical results would demonstrate statistically significant improvements in student outcomes and efficient human-AI collaboration, thereby justifying broader institutional deployment.

3. Extended multilingual support

Future versions of the system will incorporate structured support for Mauritian Creole, particularly for interface-level explanations, while preserving legal accuracy in English and French regulatory sources.

4. Longitudinal impact studies

Further evaluation over multiple academic cycles will enable assessment of long-term impacts on student retention, administrative workload reduction, and policy compliance.

5. Advanced confidence estimation

Integrating more sophisticated uncertainty modelling techniques will further improve escalation decisions, particularly for interpretative or multi-clause regulatory queries.

6. Integration with institutional systems

Linking the chatbot with existing university systems such as LMS or student information systems will enable personalised, context-aware responses while maintaining strict privacy controls.

7. Regulatory analytics and policy feedback

Aggregated query logs will be analysed to identify recurring areas of confusion, informing improvements in regulatory drafting, student communication, and policy simplification.

8. Cross-faculty and cross-institution deployment

The modular architecture will make the system readily adaptable to the other faculties within the UoM and to similar higher education institutions, enabling comparative studies across disciplines and contexts.

5.4 Final remarks

This work proposes that AI-assisted chatbots, when carefully designed with retrieval grounding, risk awareness, and governance mechanisms, can serve as trusted institutional support tools rather than mere conversational agents. The proposed system offers a conceptual approach to improving regulatory accessibility in higher education, providing a foundation for responsible AI deployment within the UoM and beyond.

References

- Andrenucci, A., & Sneiders, E. (2005). *Automated question answering: review of the main approaches*. Sydney, NSW, Australia, IEEE, pp. 514–519.
- Ashok, M., Ramasamy, K., Snehitha, G., & Keerthi, S. R., (2021). *A Systematic Survey of Cognitive Chatbots in Personalized Learning Framework*. Chennai, India, IEEE.
- Carayannopoulos, S. (2018). Using chatbots to aid transition. *International Journal of Information and Learning Technology*, 35(2), 118–129. <https://doi.org/10.1108/IJILT-10-2017-0097>
- De La Roca, M., Chan, M. M., Garcia-Cabot, A., Garcia-Lopez, E., & Amado-Salvatierra, H. (2024). The impact of a chatbot working as an assistant in a course for supporting student learning and engagement. *Computer Applications in Engineering Education*, 32(5), e22750. <https://doi.org/10.1002/cae.22750>
- Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer. <https://doi.org/10.1007/978-3-030-30371-6>
- Doshi-Velez, F., & Kim, B. (2017). *Towards a rigorous science of interpretable machine learning*. Arxiv. <https://arxiv.org/abs/1702.08608v2>
- Duarte, N., & Vardasca, R. (2023). Literature review of accreditation systems in higher education. *Educational Sciences*, 13(6), 582. <https://doi.org/10.3390/educsci13060582>
- European Commission. (2019). *Ethics guidelines for trustworthy AI*. European Commission. <https://data.europa.eu/doi/10.2759/346720>
- Følstad, A., & Brandtæg, P. B. (2017). Chatbots and the new world of HCI. *Interactions*, 24(4), 38–42. <https://doi.org/10.1145/3085558>
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
- Horvitz, E. (1999). *Principles of mixed-initiative user interfaces* (pp. 159–166). ACM Press. <https://doi.org/10.1145/302979.303030>
- Jiang, M., Gao, K., Wu, Z., & Guo, P. (2022). The influence of academic pressure on adolescents' problem behavior: Chain mediating effects of self-control, parent-child conflict, and subjective well-being. *Frontiers in Psychology*, 13, 954330. <https://doi.org/10.3389/fpsyg.2022.954330>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38. <https://doi.org/10.1145/3571730>
- Kim, Y., & Baylor, A. L. (2016). Research-based design of pedagogical agent roles: A review, progress, and recommendations. *International Journal of Artificial Intelligence in Education*, 26, 160–169. <https://doi.org/10.1007/s40593-015-0055-y>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474. <https://doi.org/10.48550/arXiv.2005.11401>
- Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed. An argument for AI in education*. Pearson.
- Luo, Q. Z., & Hsiao-Chin, L. Y. (2023). The influence of AI-powered adaptive learning platforms on student performance in chinese classrooms. *Journal of Education*, 6(3), 1–12. <https://doi.org/10.53819/81018102t4181>
- Nurshatayeva, A., Page, L. C., White, C. C., & Gehlbach, H. (2021). Are artificially intelligent conversational chatbots uniformly effective in reducing summer melt? Evidence from a randomized controlled trial. *Research in Higher Education*, 62(3), 392–402. <https://doi.org/10.1007/s11162-021-09633-z>
- Okonkwo, C. W., & Ade-Ibijola, A. (2021). Chatbots applications in education: A systematic review. *Computers and Education: Artificial Intelligence*, 2, 100033. <https://doi.org/10.1016/j.caeai.2021.100033>
- Power, D. J. (2002). *Decision support systems: Concepts and resources for managers*. Quorum Books.
- Sharda, R., Delen, D., & Turban, E. (2021). *Analytics, data science and artificial intelligence systems for decision support*. 11 ed. Pearson.
- Unesco. (2021). *AI and education: Guidance for policy-makers*. Unesco. <https://doi.org/10.54675/PCSP7350>
- Winkler, R., & Soellner, M. (2018). Unleashing the potential of chatbots in education: A state-of-the-art analysis. *Academy of Management Proceedings*, 2018, Article 15903. <https://doi.org/10.5465/ambpp.2018.15903abstract>
- Yin, J. (2021). Mindset and individual learning in the workplace: A systematic review and future agenda. *Economic Studies*, 71(1), 7–26.

To cite this article:

Beeharry, Y., Ragen, A. K., & Khoodaruth, A. (2026). Concept of an AI-assisted Regulations and Student Support Chatbot for the Faculty of Engineering at the University of Mauritius. *Artificial Intelligence Advances in Education*, 1(2), 1–17. <https://doi.org/10.66818/aiaie.v1i2.911>

Submitted: 13 January 2026

Accepted: 31 March 2026

Published: 22 May 2026

Copyright:

© 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution-NoDerivatives 4.0 International License (CC BY-ND 4.0), which permits to copy and distribute the material in any medium or format only in an unadapted form, as long as the author is named. The license allows commercial use. See <https://creativecommons.org/licenses/by-nd/4.0/>.

Artificial Intelligence Advances in Education is a peer-reviewed open access journal published by SCS Journals.

