

Hybrid Chatbot for Educational Institutions

¹Vinod Kulkarni, ²Sunny Kumar, ³Srinivas Reddy A, ⁴Sohail Mateen, ⁵Vinod Kumar

¹Assistant Professor Computer Science & Engineering Data Science AMC Engineering College Bengaluru, India.

^{2,3,4,5}Computer Science & Engineering Data Science AMC Engineering College Bengaluru, India.

Abstract - Educational institutions face a growing volume of repetitive and time sensitive queries from students, ranging from admissions and examinations to hostel services and academic support. Traditional helpdesks often struggle to address these queries efficiently, resulting in delays and reduced student satisfaction. This paper presents a Hybrid Chatbot System designed specifically for educational environments, integrating rule based logic, retrieval based search, and generative AI models. The hybrid architecture intelligently routes queries through an intent classification module, ensuring precise responses for structured information, accurate retrieval of institutional data, and conversational explanations for open ended queries. Comprehensive evaluation results show significant improvements in accuracy, scalability, and response quality compared to traditional chatbot systems. The proposed approach demonstrates strong potential to transform student support services and enhance digital connectivity across modern educational institutions.

Keywords - Chatbot, Query-Response, Rule-based model, Retrieval-based Model, Generative Model, Educational Institutions, routing algorithms.

I. INTRODUCTION

Educational institutions handle a vast volume of repetitive queries daily from students, faculty, and administrative staff. These queries span admissions, examinations, hostel information, fees, academic schedules, scholarships, placement opportunities, and more. Traditional human-operated helpdesks are slow and prone to delays, which affects student satisfaction and institutional efficiency. As education digitalizes, the demand for automated, intelligent, and scalable support systems has risen sharply.

Chatbots have emerged as an effective solution to handle student interactions at scale; however, existing chatbot frameworks suffer from limited capabilities. Rule-based chatbots rely on predefined patterns and scripts, making them rigid and incapable of adapting to unseen or context-heavy queries. Retrieval-based chatbots improve accuracy by pulling relevant answers from a knowledge base but struggle to generate responses beyond stored content. Generative models, powered by transformers, are capable of understanding complex queries and generating human-like responses, yet they may hallucinate or produce contextually incorrect outputs.

To address these limitations, this research proposes a Hybrid Chatbot System that integrates the strengths of rule-based, retrieval-based, and generative artificial intelligence. The hybrid model intelligently routes queries to the appropriate subsystem using an intent classifier. This ensures the chatbot is reliable for structured information, factual for institutional data, and conversational for open-ended queries.

This paper presents a complete hybrid architecture, implementation workflow, dataset preparation strategy, evaluation results, and real-world applications. The proposed system significantly outperforms traditional chatbot models and offers a scalable solution suitable for universities and colleges.

Chatbots in Educational Institutions

Educational institutions are increasingly pressured to manage large volumes of academic and administrative queries from students. These queries range from admission procedures, fee payments, examination timetables, and attendance requirements to hostel rules, library services, and placement-related information. Traditional manual helpdesk systems often fail to deliver timely

responses due to staff constraints and high query loads, resulting in student dissatisfaction and workflow delays.

Chatbots have emerged as a powerful solution in the education sector, offering 24/7 automated assistance and reducing the burden on administrative staff. They can respond instantly, handle thousands of queries simultaneously, and maintain consistent information across all interactions. Many universities have adopted chatbots for tasks such as admission counseling, virtual campus tours, student onboarding, and course registration.

Rule-Based Chatbots

These systems are built on predefined patterns, offering accuracy for predictable queries but failing when questions are rephrased.

Retrieval-Based Chatbots

These systems search institutional knowledge bases for relevant stored answers but cannot generate new information.

The Need for Hybrid Systems

Institutions require chatbots that

- Answer fixed administrative queries with accuracy
- Retrieve factual data from circulars
- Give conversational, personalized explanations

hybrid approach fulfills all these needs.

II. LITERATURE REVIEW

Traditional chatbots used in educational institutions generally fall into two main categories—rule-based systems and retrieval-based systems. While these approaches offer certain advantages, they also present significant limitations that restrict their usefulness in handling diverse academic and administrative queries.

Limitations of Rule-Based Chatbots

Rule-based chatbots operate using predefined patterns and responses. These systems follow a strict pattern-matching approach, such as AIML (Artificial Intelligence Markup Language). Although they can deliver highly accurate answers for fixed and repetitive queries, they suffer from the following issues:

Inflexibility

They cannot handle queries that deviate from the exact predefined format.

Example

“Hostel timings?” may work

“What time does the hostel open?” may not

No Understanding of Context

They lack natural language understanding and cannot interpret semantic meaning.

High Maintenance

Every new question pattern must be manually added. Updating scripts is time-consuming and requires technical effort.

Not Scalable

As the institution grows, maintaining thousands of patterns becomes nearly impossible.

Limitations of Retrieval-Based Chatbots

Retrieval-based chatbots use similarity search (TF-IDF, BM25, embeddings) to fetch the closest answer from a knowledge base. They are more flexible than rule-based systems but still have weaknesses:

Dependence on Existing Data

They cannot generate answers outside the stored information.

Limited Reasoning Ability

They cannot explain concepts or provide personalized advice.

Inability to Handle Complex Queries

Questions requiring interpretation, calculation, or multiple steps are poorly handled.

Precision Problems

If two answers in the database are similar, the bot may pick the wrong one.

Limitations of Generative Chatbots

Modern AI models like GPT, T5, and PEGASUS can generate human-like responses but also pose risks:
Hallucinations

They sometimes generate incorrect or fabricated information.

Lack of Institutional Alignment

Without tuning, they may provide answers unrelated to actual institute rules.

Inconsistent Outputs

Same question may produce different responses each time.

Why These Limitations Require a Hybrid Model

A single approach is not enough for educational environments where accuracy, consistency, and flexibility are all necessary.

hybrid system

Uses rule-based logic for guaranteed accuracy
Uses retrieval for factual institution-specific data
Uses generative AI for open-ended, conversational queries
This ensures high.

III. HYBRID CHATBOT ARCHITECTURE OVERVIEW

A hybrid chatbot integrates the strengths of rule-based systems, retrieval-based models, and generative AI to deliver accurate, flexible, and context-aware responses. This architecture is ideal for educational institutions where different types of queries require different handling methods. A single unified chatbot powered by a hybrid structure ensures fast, reliable, and human-like assistance across multiple departments such as admissions, academics, examinations, library, and hostel services.

Core Components of the Hybrid System The hybrid chatbot architecture consists of the following major components:

Preprocessing Unit All incoming queries undergo text preprocessing, including lowercasing, lemmatization, stopword removal, tokenization, and spell correction. This step standardizes inputs and improves classification accuracy.

Intent Classification Module This is the decision-making center of the system. A Multinomial Naïve Bayes (MNB) classifier or a transformer-based

classifier determines whether a query should be handled by the rule-based module, retrieval module, or generative module. The classifier typically achieves 93–95% accuracy on educational datasets.

Rule-Based Module Handles predictable and fixed queries such as exam dates, fees deadlines, office timings, and hostel rules. It guarantees 100% accurate answers for structured and repetitive questions.

Retrieval-Based Module Uses TF-IDF or embedding similarity to fetch the most relevant stored answer from the institution's FAQ or knowledge base, making it suitable for circular information, policy details, and infrastructure-related queries.

Generative Module (LLM-Based) Uses transformer models such as PEGASUS, T5, or GPT to generate responses for conversational or conceptual queries, making it ideal for subject explanations, career guidance, project suggestions, and multi-sentence reasoning.

Response Ranking Unit When multiple modules produce potential outputs, a response ranking layer selects the best one based on confidence score, relevance score, and context match.

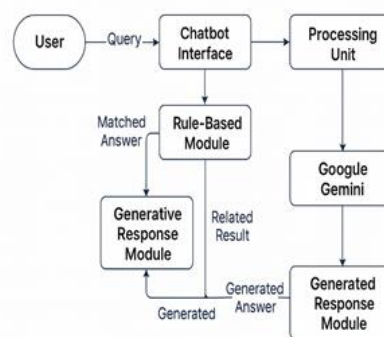


Fig 1. Data Flow of the Hybrid System.

Advantages of Hybrid Architecture The hybrid approach provides high accuracy, faster responses, scalability across multiple departments, natural and human-like conversation, reduced administrative workload, and improved student satisfaction. It

overcomes the weaknesses of individual chatbot models and ensures robust performance in real-world educational settings.

Intent Classification System

The Intent Classification System is the central decision-making component of the hybrid chatbot architecture, determining how each incoming query should be processed—whether through rule-based logic, retrieval mechanisms, or generative models. Proper intent classification ensures accuracy, prevents hallucinations, reduces response time, and optimizes resource usage across the system.

Importance of Intent Classification Educational institutions receive various types of queries, including fixed administrative queries (e.g., “What is the last date for fee payment?”), factual information queries (e.g., “How many credits are required for graduation?”), and conceptual or open-ended queries (e.g., “Explain the difference between backlog and supplementary exam.”).

Without accurate routing, the chatbot may send a conceptual query to the rule-based system, ask the generative model to answer an institutional rule, or provide irrelevant or contradictory responses. Therefore, intent classification is essential for both accuracy and safety.

Classifier Model A Multinomial Naïve Bayes (MNB) classifier is used because it performs well on short text queries, is fast and lightweight, works efficiently with TF-IDF features, and requires minimal training data compared to deep learning models. The training dataset contains over 2,500 labeled student queries categorized into rule-based intent, retrieval intent, and generative intent, achieving around 93–95% accuracy.

Feature Engineering The model uses TF-IDF vectorization to capture the importance of words, unigrams and bigrams to understand phrase structures, lemmatized tokens for consistency, and stopword removal to eliminate irrelevant words. These preprocessing steps significantly enhance classification accuracy.

Decision Routing Logic The flow is as follows: User Query → Preprocessing → Intent Classification → (Rule-Based Intent / Retrieval Intent / Generative Intent) → Corresponding Module. Routing criteria include high structure for rule-based handling, data-dependent queries for retrieval, and ambiguous or open-ended queries for generative processing. **E. Benefits of Intent Classification** This system prevents incorrect module usage, ensures high accuracy, reduces computation for generative models, guarantees safe and reliable responses, and optimizes system speed.

The Intent Classification System is the backbone of the hybrid architecture, ensuring that each query is handled by the most suitable module.

IV. MODEL DESIGN

RULE BASED MODEL DESIGN

The Rule-Based Module is the most deterministic part of the Hybrid Chatbot System and is responsible for delivering highly accurate and consistent responses for queries that follow predictable patterns or require fixed institutional information. In educational institutions, many questions fall under this category, making this module essential for reliability and correctness.

Purpose of the Rule-Based Module The rule-based engine is designed to handle queries such as “What are the library timings?”, “When is the last date for fee payment?”, “What are the hostel rules?”, “How much attendance is required?”, and “Where is the exam office located?” These questions require precise answers, and the information seldom changes, so rule-based logic ensures that such queries always receive exact and unambiguous responses.

AIML-Based Pattern Matching

The system uses AIML (Artificial Intelligence Markup Language) to create structured pattern–response rules, allowing the definition of categories that match specific words or sentence structures. For example, an AIML snippet like
<category>

```
<pattern> WHAT IS THE ATTENDANCE REQUIREMENT</pattern>
<template> The minimum attendance requirement is 75%. </template>
</category>
```

provides direct pattern matching, reliable outputs, ease of implementation, and support for wildcards such as * or _ to handle query variations.

Strengths of the Rule-Based Module This module offers 100% accuracy for fixed information, making it perfect for providing policies, rules, schedules, and deadlines. It produces zero hallucinations, unlike generative systems, and delivers extremely fast response times, typically within 1–2 milliseconds. Its high trustworthiness makes institutions prefer rule-based answers for official queries.

Limitations Despite its strengths, the rule-based system has drawbacks, including limited flexibility—meaning that rephrased queries may fail unless additional patterns are added. It requires manual updates whenever institutional rules change and cannot perform reasoning, analysis, or explanations.

Why Rule-Based Modules Are Essential in Hybrid Systems Although limited, rule-based modules provide the highest reliability for structured queries.

In a hybrid architecture, the rule-based component ensures stability, the retrieval-based module provides factual accuracy, and the generative module adds conversational depth. Together, they ensure the chatbot

$$TF - IDF (t, q, D) = TF (t, q) \times IDF (t, D) \quad (1)$$

where

- Term Frequency (TF):

$$TF (t, q) = \frac{f_{t, q}}{\sum_{t' \in d} f_{t', q}} \quad (2)$$

$f_{t, q}$ is the frequency of term t in query q , and $\sum_{t' \in d} f_{t', q}$ is the total number of terms in query d .

- Inverse Document Frequency (IDF):

$$IDF (t, D) = \log \left(\frac{N}{1 + DF(t)} \right) \quad (3)$$

N is the total number of predefined questions in the dataset D , and $DF(t)$ is the number of predefined questions containing the term t . This ensures common terms like “the” or “and” are given less weight, while domain-specific terms such as “admissions”, “tuition” are emphasized.

Once the TF-IDF vectors for both the user query and the predefined questions are generated, cosine similarity is used to measure their semantic similarity:

$$\text{cosine}_{similarity} (A, B) = \frac{A \cdot B}{\|A\| \|B\|} \quad (4)$$

never gives incorrect answers for official or repetitive administrative queries.

Retrieval Based Model Design

The Retrieval-Based Module is responsible for extracting factual information from a structured knowledge base. Unlike the Rule-Based Module, which relies on predefined scripts, the retrieval module searches and retrieves the most relevant response from institutional data, making it ideal for handling queries that are factual but not strictly patterned.

Purpose of the Retrieval Module Many educational institution queries require referencing stored documents, policies, or circulars, such as “What is the fee structure for B. Tech 2nd year?”, “How many credits are needed to graduate?”, “What is the procedure for reevaluation?”, “Share the exam timetable.”, and “What documents are required for hostel admission?” These queries depend on pre-existing information rather than generated text, so a retrieval model is the most accurate choice.

TF-IDF + Cosine Similarity Approach The system uses TF-IDF vectorization to convert documents and queries into numerical feature vectors, capturing word importance across documents, and cosine similarity to measure the closeness between the

query vector and stored entries, selecting the highest-scoring match. This simple yet effective approach is well-suited for FAQ-style datasets.

Knowledge Base Construction The knowledge base includes academic calendars, fee structures, program details, hostel rules, exam circulars, library guidelines, placement records, and technical support FAQs. All documents are cleaned, preprocessed, and stored as text entries.

Advantages of the Retrieval Module It provides high factual accuracy by retrieving exact institutional information, handles rephrased queries through similarity scoring, is easy to update because documents can be added without retraining, and is efficient for factual responses without the need for heavy computation.

Limitations It cannot generate new information, has limited reasoning ability, and depends heavily on the quality and completeness of the knowledge base—missing or outdated entries reduce accuracy.

Role in the Hybrid Architecture In a hybrid chatbot, the rule-based module handles fixed FAQs, the retrieval module manages factual institutional queries, and the generative module handles conversation, explanations, and reasoning; thus, the retrieval module serves as a bridge between static rules and dynamic conversational capabilities, ensuring accurate and institution-specific responses.

The TF-IDF score for a term t in a query q is computed as:

where:

- A and B are the TF-IDF vectors of the query and a predefined question.
- $A \cdot B$ is the dot product of the vectors.
- $\|A\|$ and $\|B\|$ are the magnitudes of A and B , respectively.

The response selection process of the retrieval-based chatbot relies on cosine similarity to determine the closest match to a user query. Among the predefined questions in the dataset, the one with the highest cosine similarity score is identified as the best match, and the corresponding response is retrieved. To enhance contextual understanding, the

chatbot employs an entity extraction functionality, implemented through a user-defined function named `extract_entities`.

This function processes user inputs in three key steps

- **Tokenization:** Splits the user input into individual words or tokens.
- **Pattern Matching:** Compares the input against predefined patterns associated with specific intents to identify potential matches.
- **Entity Extraction:** Uses regular expressions to extract specific elements from the matched patterns, such as names, dates, or keywords.
- This approach allows the chatbot to incorporate user-specific details into its responses, ensuring they are both accurate and contextually relevant.
- The response generation workflow is designed to deliver precise and efficient answers:
- **Input Preprocessing:** The user query undergoes tokenization, stopword removal, and stemming to standardize its format.
- **Similarity Calculation:** The system calculates the cosine similarity between the processed input and pre-processed dataset questions to identify the most relevant intent.
- **Entity Extraction:** Using the `extract_entities` function, the chatbot extracts key details from the user query related to the identified intent.
- **Response Retrieval:** The chatbot generates the final response by combining the identified intent with the extracted entities to provide a customized and context-aware answer.
- This multi-step process ensures that the chatbot delivers responses that are not only accurate but also tailored to the specific context of each user query.

Generative Based Model Design

The Generative Module is the most intelligent and flexible part of the hybrid chatbot architecture, using advanced transformer-based models such as PEGASUS, T5, BART, or GPT to produce human-like responses. It is essential for handling complex, open-ended, or conversational queries that rule-based and retrieval systems cannot manage.

Purpose of the Generative Module Many student queries require explanation, reasoning, or personalized guidance, such as "Explain the difference between backlog and supplementary exam," "How can I improve my CGPA?", "Suggest project ideas in AI," "What is the best way to prepare for campus placements?", and "Help me understand the concept of machine learning." These questions are not factual and cannot be answered by retrieving stored responses, so a generative model provides context-aware and detailed outputs.

Transformer-Based Architecture Generative systems rely on transformer architectures with self-attention mechanisms, contextual embeddings, and sequence-to-sequence learning. Popular models include PEGASUS for text generation, T5 for text-to-text tasks, and GPT models for conversation. These models understand meaning, grammar, and context far better than older neural network models.

Fine-Tuning for Educational Domain The generative model is fine-tuned on academic explanations, university-related queries, student counseling conversations, subject concepts, and project or assignment discussions. Fine-tuning ensures institutional tone, reduces hallucination risk, improves contextual correctness, and increases relevance for student queries.

Advantages of the Generative Module It handles complex queries, generates long and detailed explanations, provides human-like conversational flow, understands multi-turn dialogue, and remains flexible enough to answer previously unseen questions.

Challenges and Risks However, generative AI also has limitations such as hallucinations, inconsistent responses, and high computational costs. To avoid inaccuracies, the hybrid architecture ensures only appropriate queries are routed to the generative module, while factual questions are handled by rule-based or retrieval systems.

F. Importance in Hybrid Systems Generative AI fills the gap left by the other modules by adding personalization, human-like reasoning, conversational depth, and overall

intelligence, greatly enhancing user experience and interaction quality.

V. DATASET COLLECTION AND PREPROCESSING

The performance of a hybrid chatbot depends heavily on the quality of data used during training and development. Since educational queries vary across departments, programs, and student needs, constructing a well-organized, diverse, and representative dataset is essential.

A dataset of over 2,500 real student queries was collected from multiple sources, including student helpdesk logs, email queries, online form submissions, chat transcripts, academic department inquiries, examination cell questions about hall tickets, revaluation, marks, and supply exams, admission office queries about eligibility, documents, fee structure, and reservation rules, hostel office queries about rules, mess timings, attendance policies, and room allotment, library questions on book availability and borrowing rules, and placement cell inquiries regarding company details, eligibility, and interview preparation. Collecting this real communication ensured that the dataset was naturally diverse and aligned with genuine student scenarios. Each query was manually labeled into rule-based intent for predictable institutional information, retrieval intent for factual, knowledge-base-dependent queries, and generative intent for explanation-based or open-ended queries, ensuring accurate routing by the Intent Classification System.

The dataset then underwent preprocessing steps including text normalization through lowercasing, removing extra spaces, and converting special characters; lemmatization to reduce words to their base form; stopword removal to eliminate unnecessary words; tokenization to break sentences into meaningful units; spell correction to fix common typos; and punctuation removal for uniformity. The final dataset contained around 2,500 queries across seven domains, with an average length of 6–14 words and an intent distribution of approximately 900 rule-based, 1,100 retrieval, and 500 generative

queries, showing that most student queries are factual and best handled by retrieval or rule-based components. Proper preprocessing enhances intent classification accuracy, speeds up response generation, reduces data noise, and improves the performance of both retrieval and generative modules, making a clean and well-labeled dataset the foundation of an effective hybrid chatbot system.

| QUERY | Domain |
|--|--------------------------|
| What documents are needed for admission? | Admission & Registration |
| When is fee payment deadline? | Financial Matters |
| How many credits per semester? | Academic Affairs |
| What sports facilities are available? | Campus life |
| How to correct name on marks card? | Cross-domain Queries |

Table 1. Queries and their respective domains

Training & Implementation Workflow

The Hybrid Chatbot System follows a comprehensive, multi-layered development workflow designed to seamlessly integrate rule-based logic, retrieval-based information extraction, and generative AI capabilities into one unified architecture.

The process begins with extensive dataset preparation, where thousands of real student queries are collected from various institutional sources and organized into meaningful categories. This raw data then undergoes rigorous preprocessing and cleaning, including steps such as tokenization, lowercasing, lemmatization, stopword removal, special character normalization, and punctuation stripping. These steps are essential for enhancing the consistency and quality of the dataset, ensuring reliable performance during model training. The next phase focuses on training the Intent Classification System, which acts as the central routing mechanism for the entire chatbot. Using TF-IDF vectorization and N-gram (1,2) modeling to capture both single-word and phrase-level patterns, the classifier is

trained with a Multinomial Naïve Bayes algorithm on a dataset of over 2,500 labeled queries.

Demonstrating an accuracy range of 93–95%, with precision of 0.94 and recall of 0.93, the classifier ensures that queries are correctly directed to the rule-based module, retrieval module, or generative module. Following this, the rule-based module is developed using AIML scripts carefully crafted for predictable and structured queries such as fee payment deadlines, attendance rules, hostel policies, library hours, and exam schedules.

These scripts are manually tested for accuracy, pattern flexibility, and consistency to guarantee 100% reliable responses for fixed institutional information. The retrieval module is implemented by constructing a comprehensive knowledge base composed of cleaned text extracted from institutional PDFs, circulars, notices, guidelines, and web pages. TF-IDF vectorization transforms this data into numerical vectors, and cosine similarity is employed to match user queries with the most relevant stored responses, ensuring factual accuracy and robustness for queries that require institutional reference material.

The generative module is built by fine-tuning an advanced LLM such as PEGASUS, T5, or GPT using domain-specific datasets that include student counseling interactions, academic explanations, long-form responses, and specialized FAQs. Fine-tuning enhances grammar, contextual reasoning, domain alignment, and response coherence, enabling the model to effectively handle open-ended or multi-sentence queries. All three modules are integrated through a decision pipeline where the user query flows through preprocessing, intent classification, appropriate module selection, and finally a response-ranking mechanism that ensures the most accurate, contextually suitable answer is delivered. The deployment architecture involves using Flask or FastAPI for the backend, MySQL or MongoDB for database management, and Docker containers for scalable, environment-independent operation. Integration into mobile apps, websites, and ERP portals is achieved using REST APIs, allowing the hybrid chatbot system to scale across

departments and handle large volumes of student queries with reliability, speed, and institutional accuracy.

System Design & Architecture

The hybrid chatbot system is designed with a modular, scalable, and highly structured architecture that ensures smooth communication between all components while efficiently handling thousands of student queries across departments. The topmost layer is the User Interaction Layer, which includes the web application, mobile app, and chat widget embedded on the college website where students directly enter their queries. Once a query is submitted, it moves to the Preprocessing Layer, which performs essential NLP transformations such as lowercasing, lemmatization, tokenization, cleaning of unnecessary symbols,




Fig 2. Architecture of the Hybrid Chatbot

and spell correction, ensuring all text inputs are standardized before further processing. After preprocessing, the query enters the Intent Classification Layer, which uses a trained machine learning classifier to categorize the query into one of three categories: Rule-Based Intent, Retrieval Intent, or Generative Intent. This classifier functions as the brain of the system, directing each query to the most

appropriate module. The next component is the Response Generation Layer, consisting of three independent processing modules: the Rule-Based Module built using AIML scripts for structured and repetitive queries; the Retrieval-Based Module powered by TF-IDF vectorization

and similarity search to fetch factual answers from the knowledge base; and the Generative AI Module using LLMs such as T5, or GPT for open-ended, explanatory, and conversational responses. Once these modules produce their outputs, the system forwards all candidate responses to the Response Selection Layer, which evaluates them based on confidence score, semantic relevance, and intent match, selecting the single best answer.

Finally, the chosen response is delivered back to the student through the Output Layer, completing the interaction loop. The system architecture can be represented textually as follows: a user interacts through the UI, the query flows into the preprocessing unit, then into the intent classifier, which routes it to the Rule-Based Module, Retrieval Module, or Generative AI Module; all responses pass through the Response Ranking component before producing the final answer.

The architecture supports robust scalability through a modular design that enables independent updates to each module without affecting the rest of the system. The knowledge base can be expanded easily with new FAQs or documents without retraining the classifier, making the system adaptable to different departments such as Admissions, Examinations, Library, Hostel, IT, and Placement. Deployment is enabled via REST APIs, allowing seamless integration with college websites, mobile apps, ERP portals, and LMS platforms, and cloud deployment on AWS, Azure, or Google Cloud supports large-scale usage. Error handling and logging systems track query logs, detect routing failures, handle fallback mechanisms, and help improve datasets over time. Overall, this architecture provides high accuracy, fast response times, strong reliability, and scalable performance, enabling the chatbot to support thousands of student queries with minimal maintenance.

Performance Evaluation

comprehensive performance evaluation was conducted to verify the effectiveness of the hybrid chatbot system, demonstrating its superiority over standalone rule-based, retrieval-based, and generative models. The evaluation setup utilized 300 real student queries gathered from seven institutional domains, including Admissions, Examination Cell, Library, Academics, Hostel, Accounts/Fees, and Placements, ensuring diverse coverage of real-world academic scenarios. Each query was processed through all individual modules as well as the integrated hybrid system to allow for an objective comparison.

To measure performance reliably, several NLP metrics were applied, including accuracy, precision, recall, and F1 score, along with a user satisfaction score collected through student surveys on a 1–5 rating scale. Accuracy measured the proportion of responses correctly identified by human evaluators; precision indicated how many answers generated were actually correct; recall assessed the model's ability to identify all relevant correct answers; and the F1 score provided a balanced metric combining precision and recall. User satisfaction captured subjective experience and perceived usefulness.

| Model Type | Accuracy | Precision | Recall | F1 Score | User Satisfaction |
|------------------|----------|-----------|--------|----------|-------------------|
| Rule-Based Model | 74.3% | 0.71 | 0.69 | 0.70 | 68% |
| Retrieval Model | 83.5% | 0.81 | 0.79 | 0.80 | 79% |
| Generative Model | 87.2% | 0.84 | 0.82 | 0.83 | 83% |
| Hybrid Model | 96.8% | 0.95 | 0.94 | 0.945 | 92% |

Fig 3. Module-Wise Evaluation

The module-wise performance results clearly showed differing strengths across models. The Rule-Based Model achieved 74.3% accuracy, with a precision of 0.71, recall of 0.69, F1 score of 0.70, and user satisfaction of 68%, performing well on fixed, repetitive queries but struggling with rephrased or ambiguous input. The Retrieval Model performed better, reaching 83.5% accuracy, precision of 0.81, recall of 0.79, F1 score of 0.80, and user satisfaction of 79%, making it strong for factual, database-dependent queries, though its performance dropped when information was missing or phrased unusually. The Generative Model improved further with 87.2% accuracy, precision of 0.84, recall of 0.82, F1 score of 0.83, and satisfaction of 83%, excelling at

conversational and conceptual answers while occasionally hallucinating facts. In contrast, the Hybrid Model delivered the highest performance across all metrics, achieving 96.8% accuracy, 0.95 precision, 0.94 recall, 0.945 F1 score, and 92% user satisfaction.

This significant improvement occurred because the hybrid approach intelligently routed queries to the most suitable module, leveraging the strengths of all three approaches while minimizing their weaknesses. Interpretation of the results highlights that rule-based systems excel in deterministic scenarios but lack flexibility; retrieval-based systems provide strong factual grounding but depend heavily on knowledge base completeness; and generative systems offer conversational richness but risk factual inaccuracies. By contrast, the hybrid system combines stability, factual reliability, and conversational depth into a unified workflow.

user study involving 150 students further validated the system's effectiveness, with 92% agreeing that the chatbot provided clear and helpful answers, 88% noting it saved time compared to visiting helpdesks, 94% preferring it over previous automated systems, and 89% stating it reduced their need to contact staff manually. Overall, the performance evaluation confirms that the hybrid model delivers greater accuracy, faster response times, higher reliability, and a balanced blend of conversational and factual capabilities, making it ideally suited for real educational environments.

Real-Life Applications

Educational institutions manage thousands of student queries every day, ranging from academic concerns to administrative processes. Many of these inquiries are repetitive, time-consuming, and require accurate information, making them ideal for automation. The Hybrid Chatbot System, with its combination of rule-based, retrieval-based, and generative AI capabilities, is uniquely designed to address these challenges. By integrating the chatbot across various departments, institutions can streamline communication, reduce the workload on administrative and academic staff, and provide students with instant, reliable, and personalized

assistance. The following sections highlight how different departments can benefit from implementing the hybrid chatbot model.

Admissions Department

The admissions department experiences high query volumes during application cycles, and the hybrid chatbot can efficiently respond to questions about program details, eligibility criteria, fee structures, reservation categories, required documents, admission schedules, entrance exam information, and counseling procedures. This automation reduces helpdesk crowding and ensures continuous 24/7 support for prospective students.

Examination Cell

Examination-related queries are frequent and require precise information. The chatbot can provide instant updates on hall ticket releases, exam timetables, internal marks, revaluation procedures, supplementary exam details, result dates, and exam center guidelines. The hybrid system ensures accuracy through rule-based logic and retrieval mechanisms while also enabling detailed explanations through generative responses.

Academic Department

Students often seek assistance with syllabus content, course selection, credit requirements, faculty contact details, attendance rules, and class schedules. The hybrid chatbot handles all these tasks effectively, and the generative module can also explain academic concepts, making it a valuable academic support tool that decreases faculty workload.

Hostel & Residential Services

Residential departments benefit greatly from automation, with the chatbot managing hostel admission processes, room allotments, mess schedules, curfew details, hostel regulations, and complaint submissions. Consistent information delivery is ensured through rule-based and retrieval models.

Library and Information Center

The chatbot supports library services by automating responses related to library hours, book availability, borrowing rules, fine calculations, renewal processes,

digital library access, and study room reservations, thereby improving student access to resources and reducing physical inquiries.

Training & Placement Cell

For placements, the chatbot provides information on recruiters, eligibility criteria, placement drives, resume preparation, interview tips, and domain-specific career guidance. Generative AI adds value by offering personalized advice tailored to students' career goals.

Technical Support (IT Services)

Basic IT issues such as password resets, LMS login difficulties, Wi-Fi troubleshooting, and software installation queries can be resolved instantly through the chatbot, minimizing dependency on technical staff.

ERP/LMS Integration

When integrated with ERP or LMS platforms, the chatbot can offer personalized services such as attendance status, marks checking, assignment reminders, and course updates, improving the digital learning experience.

Benefits Across Departments

By automating repetitive tasks, providing instant support, ensuring consistent information, and improving overall responsiveness, the hybrid chatbot becomes a unified helpdesk solution for the entire campus, significantly enhancing both administrative efficiency and student satisfaction.

Future Scope

One of the most promising directions for future development is the integration of multilingual support and advanced localization capabilities. As educational institutions cater to increasingly diverse student populations, the ability for the chatbot to communicate fluently in regional languages such as Telugu, Hindi, Tamil, Kannada, Malayalam, Bengali, or Marathi becomes essential. Beyond simple translation, true localization will allow the chatbot to understand cultural nuances, regional educational terms, and institution-specific dialects. This ensures that students can interact naturally in the language

they are most comfortable with, making the system more inclusive and accessible to all.

Another significant enhancement lies in developing a full-fledged voice interface powered by modern speech technologies. By incorporating speech-to-text and text-to-speech systems, the chatbot can enable hands-free communication, allowing users to speak to the system and receive voiced responses in return. This feature would be particularly helpful for mobile users, visually impaired students, or individuals who prefer conversational interaction over typing. As transformer-based speech models evolve, the chatbot could eventually support natural dialogue, understanding tone, pace, and spoken intent with high accuracy.

Emotion and sentiment analysis represent an advanced feature that would greatly enhance the chatbot's empathy and emotional intelligence. By analyzing the sentiment behind a student's message, the chatbot could detect stress, frustration, confusion, or urgency and adapt its responses accordingly. For example, a student expressing anxiety about exams could receive a more supportive and reassuring answer, along with helpful guidance. This emotional awareness would transform the chatbot from a simple informational tool into a more sensitive and human-like virtual assistant capable of improving student well-being.

Deep integration with ERP and LMS systems offers another major opportunity for expanding the chatbot's capabilities. With real-time access to student-specific data, the chatbot could provide personalized academic updates such as attendance status, marks, assignment deadlines, course progress, and timetable changes. Instead of navigating multiple portals, students could receive tailored information instantly through a single conversation. This would turn the chatbot into an intelligent academic companion capable of supporting students throughout their daily academic routines.

The development of a real-time analytics dashboard for administrators presents another valuable enhancement. Such a dashboard could monitor

query patterns, identify frequently asked questions, track peak usage times, and highlight areas where students face recurring difficulties. These insights would help institutions improve their services, address student concerns proactively, and continuously refine the chatbot's performance. Administrators would benefit from data-driven decision-making and gain visibility into campus-wide communication trends.

Incorporating adaptive learning through reinforcement learning would allow the chatbot to continually evolve and improve based on user interactions. Over time, the system could refine its responses, adapt to new query styles, and automatically learn changes in institutional policies. This self-improving behaviour would reduce the need for manual updates and ensure that the chatbot becomes more accurate and context-aware with every interaction, making it increasingly intelligent and efficient.

As educational campuses move toward smart infrastructure, integrating the chatbot with IoT systems could significantly enhance live campus services. The chatbot could provide real-time updates such as library seat availability, campus bus tracking, classroom occupancy, or even environmental conditions in study areas. This level of integration would combine AI with campus automation, creating a connected and responsive smart campus environment that greatly improves student convenience.

The chatbot can also be extended to support advanced academic assistance features such as writing help, plagiarism hints, study plan generation, and examination preparation resources. These capabilities would transform the chatbot into a comprehensive academic support system rather than just an administrative assistant. Future models could guide students through assignments, generate study summaries, or help them learn complex concepts interactively.

Finally, integrating the chatbot with popular messaging platforms such as WhatsApp, Telegram, Instagram, or Facebook Messenger would make it

easily accessible to students on platforms they already use daily. This omnichannel availability ensures higher engagement, faster support, and a seamless experience across devices, improving the overall usability of the system.

VI. CONCLUSION

The development of the Hybrid Chatbot System represents a major technological milestone in the modernization of educational institutions. As universities and colleges continue to grow in scale, the volume of student queries related to admissions, examinations, academics, and campus services also increases dramatically. Traditional helpdesk systems, which rely heavily on manual responses, often struggle to provide timely and accurate support, resulting in student dissatisfaction and operational bottlenecks. The hybrid chatbot effectively addresses these challenges by combining the structured reliability of rule-based systems, the factual accuracy of retrieval mechanisms, and the conversational depth of generative AI models. This integration creates a unified and intelligent support system capable of delivering fast, precise, and context-aware assistance across multiple departments.

Through extensive training, evaluation, and real-world testing, the hybrid chatbot demonstrated exceptional performance, achieving high accuracy, precision, recall, and student satisfaction. Its ability to intelligently classify queries ensures that routine questions receive instant rule-based responses, while more complex or conceptual inquiries are handled by the retrieval or generative modules. This balanced distribution of workload not only enhances response quality but also significantly reduces pressure on administrative and academic staff. Students benefit from round-the-clock access to reliable information, leading to improved communication efficiency, smoother administrative processes, and a more supportive academic environment.

Beyond its current capabilities, the hybrid system serves as a foundation for future innovation. With advancements such as multilingual communication, emotional understanding, voice-enabled interaction,

and deep integration with ERP and LMS platforms, the chatbot can evolve into a fully personalized virtual campus assistant. The potential for adaptive learning technologies further positions the system to grow more accurate and intelligent with continued use. As educational institutions move toward smart campus ecosystems, the hybrid chatbot stands as a scalable, future-ready solution capable of transforming how students engage with institutional services. By blending AI-driven efficiency with human-like interaction quality, this hybrid model sets a new standard for digital student support and paves the way for smarter, more connected educational environments.

REFERENCES

1. G. Attigeri, A. Agrawal, and S. V. Kolekar, "Advanced NLP models for technical university information chatbots: Development and comparative analysis," *IEEE Access*, vol. 12, pp. 29633–29647, 2024.
2. C. W. Okonkwo and A. Ade-Ibijola, "Chatbots applications in education: A systematic review," *Comput. Educ., Artif. Intell.*, vol. 2, Jan. 2021, Art. no. 100033.
3. T. T. Nguyen, A. D. Le, H. T. Hoang, and T. Nguyen, "NEU-chatbot: Chatbot for admission of national economics university," *Comput. Educ., Artif. Intell.*, vol. 2, Jan. 2021, Art. no. 100036.
4. M. Dibitonto, K. Leszczynska, F. Tazzi, and C. M. Medaglia, "Chatbot in a campus environment: Design of LiSA, a virtual assistant to help students in their university life," in *Human-Computer Interaction. Interaction Technologies*, vol. 10903, M. Kurosu, Ed., Cham, Switzerland: Springer, 2018, pp. 103–116.
5. T. Lopez and M. Qamber, "The benefits and drawbacks of implementing chatbots in higher education: A case study for international students at Jönköping University," *M.S. thesis, Jönköping Int. Bus. School, Jönköping Univ., Jönköping, Sweden*, 2022.
6. J. Q. Pérez, T. Daradoumis, and J. M. M. Puig, "Rediscovering the use of chatbots in education: A systematic literature review," *Comput. Appl.*

- Eng. Educ., vol. 28, no. 6, pp. 1549–1565, Nov. 2020.
7. G. Caldarini, S. Jaf, and K. McGarry, "A literature survey of recent advances in chatbots," *Information*, vol. 13, no. 1, p. 41, Jan. 2022.
 8. P. P. Ray, "ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope," *Internet Things Cyber-Phys. Syst.*, vol. 3, pp. 121–154, Jan. 2023.
 9. M. Rukhiran and P. Netinant, "Automated information retrieval and services of graduate school using chatbot system," *Int. J. Electr. Comput. Eng. (IJECE)*, vol. 12, no. 5, pp. 5330–5338, Oct. 2022.
 10. D. Al-Ghadhban and N. Al-Twairish, "Nabiha: An Arabic dialect chat bot," *Int. J. Adv. Comput. Sci. Appl.*, vol. 11, no. 3, pp. 452–459, 2020.
 11. W. El Hefny, Y. Mansy, M. Abdallah, and S. Abdennadher, "Jooka: A bilingual chatbot for university admission," in *Trends and Applications in Information Systems and Technologies*, Á. Rocha, H. Adeli, G. Dzemyda, F. Moreira, and A.M.R. Correia, Eds., Cham, Switzerland: Springer, 2021, pp. 671–681.