



Original Article

Generative AI Frameworks for Digital Academic Advising and Intelligent Student Support Systems

Dilliraja Sundar¹, Yashovardhan Jayaram², Jayant Bhat³
^{1,2,3}Independent Researcher, USA.

Abstract - The growing scale and complexity of higher education have exposed limitations in traditional academic advising models, which are often manual, reactive, and constrained by limited human resources. As institutions serve increasingly diverse student populations across physical and digital learning environments, there is a critical need for intelligent, scalable, and personalized advising solutions. This paper presents a comprehensive framework for Generative Artificial Intelligence (GenAI)-driven digital academic advising and intelligent student support systems that integrates large language models, predictive analytics, and institutional academic data. The proposed framework enables natural language interaction, personalized academic planning, early risk identification, and career-oriented guidance by leveraging student records, learning management data, and policy knowledge bases. Generative AI enhances advising by providing context-aware, explainable recommendations and supporting scenario-based academic decision-making. To ensure responsible adoption, the framework incorporates ethical AI principles, including transparency, bias mitigation, privacy preservation, and human-in-the-loop governance. A modular, cloud-native, and API-driven architecture is proposed to support scalability and seamless integration with existing Student Information Systems and Learning Management Systems. Experimental insights and related studies indicate that AI-enabled advising can significantly improve student engagement, advising efficiency, and retention outcomes while reducing advisor workload. Overall, the framework demonstrates how generative AI can transform academic advising into a proactive, intelligent, and continuously learning support ecosystem that aligns institutional goals with student success and long-term educational outcomes.

Keywords - Generative Artificial Intelligence, Digital Academic Advising, Intelligent Student Support Systems, Large Language Models, Educational Data Analytics, Explainable AI, Decision Support Systems, Higher Education Technology.

1. Introduction

The increasing complexity of higher education systems, driven by expanding curricula, diverse student populations, and evolving career pathways, [1,2] has placed significant pressure on traditional academic advising models. Conventional advising practices are largely manual, advisor-centric, and reactive, relying on scheduled interactions and static rule-based guidelines. While effective at small scales, these approaches often struggle to provide timely, personalized, and continuous support to students, particularly in large universities and online or hybrid learning environments. As a result, students may experience delayed guidance, suboptimal course selections, and limited awareness of academic risks, which can negatively impact retention, progression, and overall academic success.

Recent advancements in artificial intelligence (AI) and educational data analytics have opened new possibilities for transforming academic advising into a more intelligent and student-centered process. Predictive models and learning analytics have been applied to identify at-risk students, recommend courses, and analyze academic performance trends. However, many existing systems remain narrowly focused, lack natural interaction capabilities, and depend heavily on predefined rules or historical correlations. These limitations restrict their ability to adapt to dynamic student contexts, interpret complex academic policies, and provide explainable, conversational guidance that students can easily understand and trust.

Generative Artificial Intelligence (GenAI), particularly large language models, offers a paradigm shift by enabling context-aware reasoning, natural language interaction, and proactive decision support. When integrated with institutional data and governance mechanisms, GenAI-driven academic advising systems can deliver personalized recommendations, simulate academic planning scenarios, and support both students and advisors in real time. This paper introduces a structured framework for generative AI-based digital academic advising and intelligent student support systems, highlighting its architectural components, ethical considerations, and potential to enhance student engagement, decision quality, and long-term educational outcomes.

2. Related Work and Literature Review

2.1. Digital Academic Advising Platforms

Digital academic advising has traditionally been supported by enterprise-scale Student Information Systems (SIS) and Learning Management Systems (LMS), which serve as the backbone of academic data management in higher education. [3-5] SIS platforms centralize student information such as enrollment records, grades, attendance, degree requirements, and financial aid details, enabling advisors and administrators to monitor academic progress and compliance with institutional policies. Systems such as PowerSchool, Skyward, and Anthology Student provide degree audit capabilities, progress dashboards, and reporting tools that support retention strategies and administrative decision-making. These platforms significantly reduce manual record-keeping errors and improve operational efficiency, but they primarily function as data repositories and reporting systems rather than intelligent advisory tools.

Learning Management Systems, including Blackboard and Moodle, complement SIS platforms by managing course delivery, assessments, and student–faculty interactions. LMS data captures fine-grained learning behaviors such as assignment submissions, quiz performance, and participation metrics. When integrated with SIS, these systems enable a unified view of both academic outcomes and learning engagement, allowing advisors to detect performance issues earlier. However, despite their analytical capabilities, most SIS–LMS ecosystems rely on static rules and predefined workflows, limiting their ability to provide adaptive, personalized, and conversational advising support.

2.2. AI and Machine Learning in Student Support

The application of artificial intelligence and machine learning in student support has been widely explored, particularly in the context of predictive analytics and early warning systems. Machine learning models analyze historical and real-time student data such as attendance patterns, grades, assessment scores, and socio-economic indicators to identify students at risk of poor performance or dropout. Algorithms including Naive Bayes, decision trees, random forests, and gradient boosting techniques have demonstrated strong predictive accuracy in retention and performance forecasting. Empirical studies report that data-driven interventions informed by these models can improve graduation and retention rates by approximately 15–20%.

Despite these benefits, most AI-driven advising tools focus on prediction rather than explanation or interaction. The outputs are often numerical risk scores or alerts that require human interpretation by advisors. This dependency limits scalability and reduces direct student engagement. Moreover, traditional machine learning models struggle to incorporate unstructured data, such as advising notes or policy documents, and lack the reasoning capabilities needed for complex academic planning scenarios.

2.3. Conversational Agents in Education

Conversational agents and chatbots have emerged as a complementary approach to enhance accessibility and responsiveness in academic support services. These systems provide round-the-clock assistance for common student queries related to course schedules, assessment deadlines, administrative procedures, and academic resources. In educational settings, conversational agents have also been used for virtual tutoring, formative feedback, and student engagement, with studies indicating improvements in satisfaction and learning outcomes due to immediate and personalized responses.

Domain-specific conversational agents outperform generic chatbots by leveraging contextual academic data and institutional knowledge, enabling more accurate and relevant guidance. However, existing educational chatbots are often limited to scripted interactions or narrow task domains, requiring escalation to human advisors for complex or sensitive issues. This highlights the need for more advanced generative AI–based conversational frameworks that combine natural language understanding, reasoning, and institutional data integration while maintaining human oversight and ethical governance.

3. Problem Definition and Design Requirements

3.1. Academic Advising Challenges in Large-Scale Institutions

Large-scale higher education institutions face substantial challenges in delivering effective academic advising due to the diversity and volume of their student populations. [6-8] Students differ widely in academic preparedness, learning pace, socio-economic background, career aspirations, and personal constraints, resulting in highly heterogeneous advising needs. Traditional one-size-fits-all advising models are inadequate for addressing such variability, often leading to generic guidance that fails to align with individual student contexts. Additionally, institutions must support undergraduate, postgraduate, part-time, online, and international learners simultaneously, each governed by distinct academic policies and progression rules. Another critical challenge is the growing demand for real-time and continuous guidance. Students increasingly expect immediate support for course selection, prerequisite validation, academic risk alerts, and policy clarification, particularly in digital and hybrid learning environments. However, human advisors are constrained by limited availability and high advisor-to-student ratios, making timely interventions

difficult. As a result, advising often becomes reactive rather than proactive, addressing problems only after academic difficulties arise. These challenges highlight the need for intelligent, scalable systems capable of delivering personalized, context-aware, and real-time academic advising support without increasing institutional workload.

3.2. Functional and Non-Functional System Requirements

An effective digital academic advising system must satisfy both functional and non-functional requirements to ensure reliability, usability, and institutional trust. Functionally, the system should provide accurate academic recommendations, including course planning, degree audits, prerequisite validation, and early risk identification based on comprehensive student data. It must support natural language interaction, enabling students to query academic policies and receive personalized guidance in an intuitive manner. Non-functional requirements are equally critical. Scalability is essential to support thousands of concurrent users without performance degradation, particularly during peak academic periods such as enrollment and examinations. Explainability is required to ensure that AI-generated recommendations can be understood and justified by both students and advisors, fostering trust and informed decision-making. Data privacy and security must be embedded by design, ensuring that sensitive academic and personal information is protected throughout data processing and storage. Additionally, system interoperability through standardized APIs is necessary to integrate seamlessly with existing SIS, LMS, and institutional platforms.

3.3. Ethical, Privacy, and Regulatory Constraints

The deployment of AI-driven academic advising systems must operate within strict ethical, privacy, and regulatory boundaries. Compliance with educational data protection regulations such as the Family Educational Rights and Privacy Act (FERPA) and the General Data Protection Regulation (GDPR) is mandatory, requiring explicit controls over data access, consent management, and data minimization. Student data must be used solely for legitimate educational purposes, with transparent disclosure of how recommendations are generated and utilized. Ethical considerations also include bias and fairness, as AI models trained on historical academic data may inadvertently reinforce existing inequalities related to gender, socio-economic status, or prior academic performance. Ensuring fairness requires careful dataset curation, continuous bias evaluation, and model auditing mechanisms. Furthermore, human-in-the-loop governance is essential to prevent over-reliance on automated decisions, particularly in high-stakes scenarios such as academic probation or graduation eligibility. Addressing these ethical and regulatory constraints is critical to ensuring responsible, trustworthy, and sustainable adoption of generative AI in academic advising systems.

4. Generative AI–Driven Advising Framework

4.1. Overall System Architecture

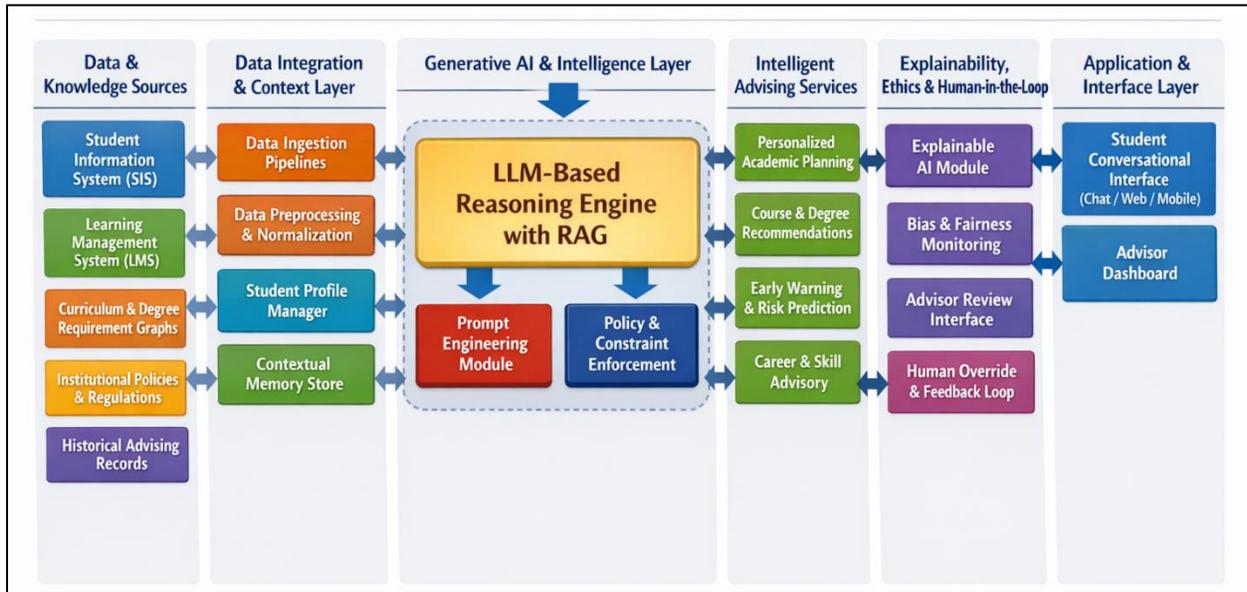


Fig 1: Overall Generative AI–Driven Academic Advising System Architecture

The figure illustrates the end-to-end architecture of the proposed generative AI–driven academic advising and intelligent student support framework. [9-11] The architecture is organized into clearly defined layers that collectively enable personalized, explainable, and scalable advising services. At the foundation, the Data and Knowledge Sources layer aggregates structured and

unstructured institutional data, including Student Information Systems (SIS), Learning Management Systems (LMS), curriculum and degree requirement graphs, institutional policies, and historical advising records. These sources provide authoritative and policy-compliant knowledge that ensures the system’s recommendations remain aligned with official academic regulations and student histories.

The Data Integration and Context layer is responsible for transforming raw institutional data into contextualized inputs suitable for AI reasoning. This layer includes data ingestion pipelines, preprocessing and normalization modules, and a student profile manager that consolidates academic progress, performance trends, and interaction history. A contextual memory store maintains short-term and long-term advising context, enabling continuity across multiple interactions. By unifying heterogeneous data into a coherent representation, this layer ensures that downstream generative reasoning is accurate, personalized, and context-aware.

At the core of the framework lies the Generative AI and Intelligence layer, centered on an LLM-based reasoning engine enhanced with Retrieval-Augmented Generation (RAG). This engine dynamically retrieves relevant institutional knowledge and combines it with generative reasoning to produce grounded and policy-consistent responses. Supporting components such as prompt engineering modules and policy constraint enforcement mechanisms control model behavior, reduce hallucination, and ensure compliance with academic rules. This design balances the flexibility of large language models with the reliability required for high-stakes academic decision support.

The upper layers of the architecture focus on service delivery, governance, and user interaction. The Intelligent Advising Services layer provides core functionalities such as personalized academic planning, course and degree recommendations, early warning and risk prediction, and career and skill advisory. These services are governed by the Explainability, Ethics, and Human-in-the-Loop layer, which introduces explainable AI modules, bias and fairness monitoring, advisor review interfaces, and human override mechanisms to maintain trust and accountability. Finally, the Application and Interface layer delivers these capabilities through student-facing conversational interfaces and advisor dashboards, enabling seamless interaction while preserving transparency and institutional control.

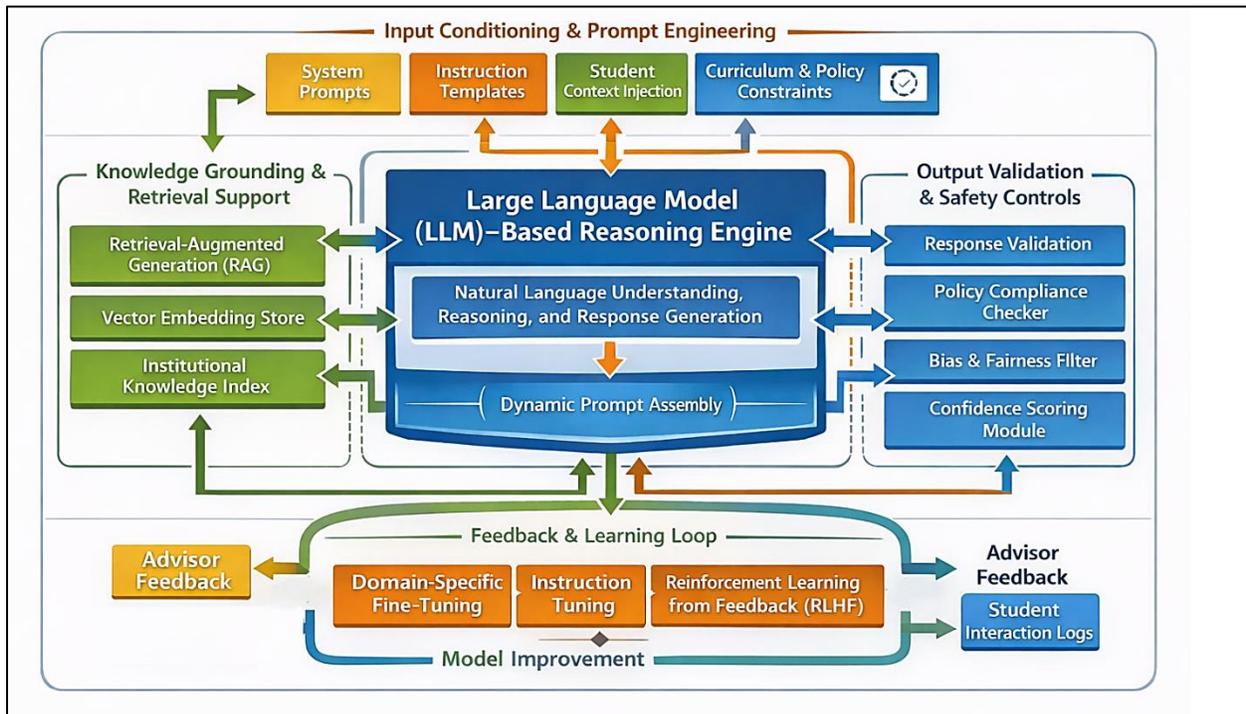


Fig 2: LLM-Based Reasoning Engine with Prompt Engineering, Knowledge Grounding, and Safety Control Mechanisms

4.2. Data Sources and Knowledge Integration

The data sources and knowledge integration layer consolidates heterogeneous academic and institutional information required for effective digital advising. Core inputs include Student Information Systems (SIS) that provide authoritative records on enrollment, grades, and degree progress, and Learning Management Systems (LMS) that capture fine-grained learning behaviors

such as assessment performance and engagement patterns. Curriculum and degree requirement graphs formally encode prerequisite structures, credit constraints, and program pathways, enabling rule-aware academic reasoning. Institutional policy documents, including academic regulations and advising guidelines, are integrated as structured knowledge artifacts to ensure compliance and consistency. Through standardized ingestion pipelines and semantic indexing, these diverse data sources are unified into a coherent knowledge base that supports accurate, policy-aligned, and context-aware AI-driven advising.

4.3. Generative Model Layer

The generative model layer serves as the intelligence core of the advising framework, leveraging large language models (LLMs) for reasoning, natural language understanding, and response generation. LLMs are employed to interpret student queries, synthesize institutional knowledge, and generate personalized advising recommendations in an explainable manner. Prompt engineering strategies are used to structure inputs, constrain outputs, and embed academic rules and context directly into model interactions. Fine-tuning and retrieval-augmented generation techniques further adapt the models to institutional domains, improving factual accuracy and reducing hallucination risks. Together, these mechanisms enable reliable, context-sensitive, and scalable generation of academic guidance while maintaining alignment with formal policies and advising objectives.

4.4. Context Management and Student Profiling

Context management and student profiling enable personalized and continuous advising experiences across multiple interactions. This component maintains dynamic student profiles that capture academic history, current progress, preferences, and prior advising exchanges. Contextual memory mechanisms preserve short-term conversational context as well as long-term advising states, allowing the system to provide consistent recommendations over time. By continuously updating profiles based on new academic data and interactions, the framework adapts guidance to evolving student needs, ensuring relevance and continuity. This context-aware design is essential for delivering personalized, longitudinal academic support that mirrors human advising relationships while operating at institutional scale.

5. Intelligent Student Support Services

5.1. Personalized Academic Planning and Course Recommendations

Personalized academic planning is a core capability of intelligent student support systems, enabling students to navigate complex degree requirements with clarity and confidence. [12-14] By leveraging generative AI models integrated with Student Information Systems and institutional policy repositories, the system can automatically generate individualized degree pathways aligned with a student's academic history, performance trends, and declared or evolving goals. These pathways dynamically adapt to changes such as course failures, credit transfers, or modified graduation timelines, offering scenario-based planning options that allow students to explore alternative academic trajectories. Prerequisite-aware guidance further enhances decision-making by validating course eligibility in real time, identifying unmet requirements, and recommending optimal sequencing to balance workload and academic rigor. Unlike static rule-based systems, AI-driven planning accounts for contextual factors such as course availability, historical success rates, and student learning preferences. Natural language interfaces allow students to query what-if scenarios, such as delaying a core course or adding a minor, and receive clear, explainable recommendations. This personalized and interactive approach reduces advising bottlenecks, minimizes enrollment errors, and empowers students to take greater ownership of their academic progress while maintaining alignment with institutional policies.

5.2. Early Warning and Risk Prediction

Early warning and risk prediction services play a critical role in improving student retention and academic success. Intelligent support systems continuously analyze academic performance, attendance records, learning management system engagement, and behavioral indicators to identify students at risk of underperformance or attrition. Machine learning and generative AI models go beyond static risk scores by contextualizing risk factors and explaining why a student may be struggling, such as inconsistent participation, declining assessment performance, or unmet prerequisites. This interpretability enables timely and targeted interventions rather than generic alerts. Proactive AI-driven interventions may include personalized study recommendations, reminders for academic milestones, referrals to tutoring or counseling services, and adaptive workload suggestions. Importantly, these interventions can be delivered directly to students through conversational interfaces while simultaneously notifying advisors for oversight. By shifting from reactive support to continuous monitoring and early engagement, institutions can address academic challenges before they escalate, reduce dropout rates, and foster a supportive learning environment that responds dynamically to individual student needs.

5.3. Career and Skill Development Advising

Career and skill development advising extends intelligent student support beyond academic success to long-term employability. Generative AI systems can perform skill gap analysis by mapping a student's completed coursework, projects, and extracurricular activities against industry-defined skill frameworks and labor market data. This analysis identifies strengths,

deficiencies, and emerging skill requirements relevant to targeted career paths. Based on these insights, the system provides industry-aligned recommendations, such as elective courses, certifications, internships, or project-based learning opportunities that enhance career readiness. Natural language interaction enables students to explore career scenarios, compare roles, and understand how academic choices influence employability outcomes. Additionally, the system can align recommendations with evolving industry trends, ensuring relevance in rapidly changing job markets. By integrating academic advising with career guidance, intelligent student support systems help students make informed, strategic decisions that bridge the gap between education and employment, supporting both institutional graduate outcomes and individual career aspirations.

6. Explainability, Trust, and Human-in-the-Loop Design

6.1. Explainable Generative AI for Academic Decisions

Explainability is a critical requirement for the adoption of generative AI in academic advising, where decisions can significantly influence student progression and outcomes. [15-17] Explainable generative AI systems provide transparent recommendation rationales by clearly articulating the factors, rules, and data sources that contribute to a given advisory output. Instead of presenting opaque suggestions, the system contextualizes recommendations by referencing degree requirements, prerequisite structures, historical performance patterns, and institutional policies in natural language. For example, when recommending a specific course sequence, the system can explain how prior grades, credit load balance, and prerequisite completion influenced the recommendation. Such transparency enhances student understanding and enables advisors to validate and justify AI-supported decisions. Explainability also supports accountability by allowing institutions to audit system behavior and detect inconsistencies or errors in reasoning. By aligning AI outputs with human-interpretable explanations, generative advising systems foster trust, reduce resistance to automation, and ensure that academic decisions remain understandable, reviewable, and defensible within formal educational governance structures.

6.2. Advisor–AI Collaboration Models

Effective academic advising systems are designed to augment, rather than replace, human advisors through well-defined advisor–AI collaboration models. In a decision support paradigm, generative AI functions as an intelligent assistant that analyzes data, generates recommendations, and highlights potential risks, while final decisions remain under human control. This model preserves professional judgment and allows advisors to incorporate qualitative factors such as student motivation, personal circumstances, and emotional well-being that AI systems may not fully capture. In contrast, full decision automation is generally unsuitable for high-stakes academic scenarios, as it risks over-reliance on algorithmic outputs and reduces institutional accountability. Human-in-the-loop design ensures that advisors can review, modify, or override AI-generated guidance, particularly in cases involving academic probation, graduation eligibility, or special accommodations. By embedding AI within existing advising workflows, institutions can improve efficiency and consistency without compromising ethical responsibility, thereby strengthening trust among students, advisors, and administrators.

6.3. Bias Detection and Fairness Assurance

Bias detection and fairness assurance are essential to prevent generative AI systems from reinforcing existing inequities in academic advising. Historical academic data may reflect systemic biases related to gender, socio-economic background, ethnicity, or prior educational access, which can influence model behavior if left unaddressed. Fairness-aware system design incorporates continuous monitoring mechanisms to detect disparate impacts across demographic groups and academic cohorts. Techniques such as bias audits, counterfactual testing, and balanced training datasets help ensure that recommendations are equitable and do not disadvantage specific populations. Additionally, explainable outputs allow advisors to identify and question potentially biased recommendations, reinforcing human oversight. Fairness assurance also involves policy-level constraints that limit the use of sensitive attributes and enforce ethical guidelines in decision-making. By integrating bias mitigation strategies with transparent governance and human review, generative AI–driven advising systems can promote inclusive, ethical, and trustworthy academic support that aligns with institutional values and regulatory expectations.

7. Implementation and Deployment Considerations

7.1. Model Training, Fine-Tuning, and Evaluation

The successful deployment of generative AI–driven academic advising systems depends heavily on effective model training, fine-tuning, and continuous evaluation. [18-20] Pre-trained large language models provide strong general reasoning and language capabilities but must be adapted to institutional contexts using domain-specific data such as academic catalogs, degree requirements, advising policies, and anonymized student records. Institutional data adaptation involves careful data preprocessing, normalization, and labeling to ensure consistency and relevance while preserving privacy. Fine-tuning techniques, including supervised instruction tuning and retrieval-augmented generation, enable models to generate accurate, policy-compliant responses grounded in authoritative institutional knowledge. Evaluation must extend beyond traditional accuracy metrics to include

explainability, fairness, and user satisfaction. Regular performance assessments using historical advising scenarios and real-world pilot deployments help identify gaps in recommendation quality and reasoning. Continuous feedback from advisors and students further supports iterative model refinement, ensuring that the system remains aligned with evolving curricula, policies, and educational objectives.

7.2. Cloud-Native and Scalable Deployment

Cloud-native deployment is essential for supporting the scalability, availability, and flexibility required by large-scale academic advising systems. A microservices-based architecture enables modular development, allowing components such as conversational interfaces, recommendation engines, analytics services, and policy reasoning modules to be independently deployed and scaled. API-driven communication ensures seamless integration with existing institutional platforms, including SIS, LMS, and identity management systems. This architectural approach supports elastic scaling during peak usage periods, such as course registration and examinations, without compromising system performance. Containerization and orchestration technologies further enhance resilience, enabling rapid updates and fault isolation. Additionally, cloud-native deployment facilitates cross-campus and multi-institutional adoption by supporting multi-tenancy and configuration-driven customization. By decoupling system components and leveraging standardized interfaces, institutions can achieve operational efficiency, faster innovation cycles, and long-term maintainability of AI-driven advising platforms.

7.3. Security and Access Control Mechanisms

Security and access control mechanisms are fundamental to protecting sensitive academic and personal data within AI-enabled advising systems. Robust identity and access management frameworks enforce role-based access control, ensuring that students, advisors, and administrators can only access authorized information and functionalities. Secure authentication protocols, such as OAuth and single sign-on, reduce unauthorized access risks while improving user experience. Data encryption at rest and in transit safeguards student records from interception and breaches. Additionally, audit logging and continuous monitoring enable institutions to track system usage, detect anomalies, and support regulatory compliance. Fine-grained access policies are particularly important in generative AI systems, where models may interact with multiple data sources. By embedding security controls at every architectural layer, institutions can ensure confidentiality, integrity, and availability of academic data while maintaining trust and compliance with educational data protection standards.

8. Experimental Evaluation and Case Study

8.1. Dataset Description and Experimental Setup

The experimental evaluation was conducted using real academic data collected from 100 undergraduate Information Technology (IT) students enrolled in second to fourth years at Ho Chi Minh City University of Education. The dataset included structured academic records such as course grades, cumulative performance history, and curriculum mappings, along with unstructured interaction logs generated through student conversations with the FIT-Advisor chatbot. Institutional regulations, curriculum descriptions, and foundational IT knowledge were synchronized from official faculty manuals to construct a verified academic knowledge base.

The proposed system was implemented using a Rasa-based generative AI conversational framework, augmented with deep learning-based grade prediction models and PhoBERT, a Vietnamese language model, for natural language understanding and semantic processing. The chatbot supported personalized course recommendations, natural language queries on academic regulations, and foundational IT concepts. The pilot deployment was conducted over a controlled evaluation period in 2024, accessible through a web-based interface. System usability and effectiveness were assessed using pre- and post-deployment student surveys, advisor feedback, and system-level performance metrics.

Table 1: Dataset Components and Sources

Dataset Component	Size / Description	Source
Student Records	100 IT students' academic history and scores	University database
Interaction Logs	Query-response pairs	FIT-Advisor chatbot sessions
Knowledge Base	Regulations, curricula, IT Q&A	Faculty manuals

8.2. Performance Metrics and Quantitative Results

System performance was evaluated across predictive accuracy, response quality, latency, and user satisfaction. The grade prediction module achieved 90% accuracy, demonstrating the effectiveness of deep learning representations combined with contextual academic features. For advising-related queries, 62% of responses were rated as completely accurate, reflecting strong

alignment between generative outputs and institutional policies. User experience surveys revealed that 71% of students rated the system as extremely useful, with an average usability score of 4.7 out of 5.

From a system efficiency perspective, the average response latency remained below 2 seconds, meeting real-time interaction requirements for digital advising platforms. These findings align with external evidence from the Tyton Partners 2024 survey (n > 3,000), which reported that AI-enabled advising significantly reduces transactional advising workload and enables more holistic student–advisor interactions.

Table 2: Key Performance Metrics

Metric	Value
Grade Prediction Accuracy	90%
Advising Response Accuracy	62% (completely accurate)
Student Satisfaction	71% Extremely useful
Response Latency	< 2 seconds

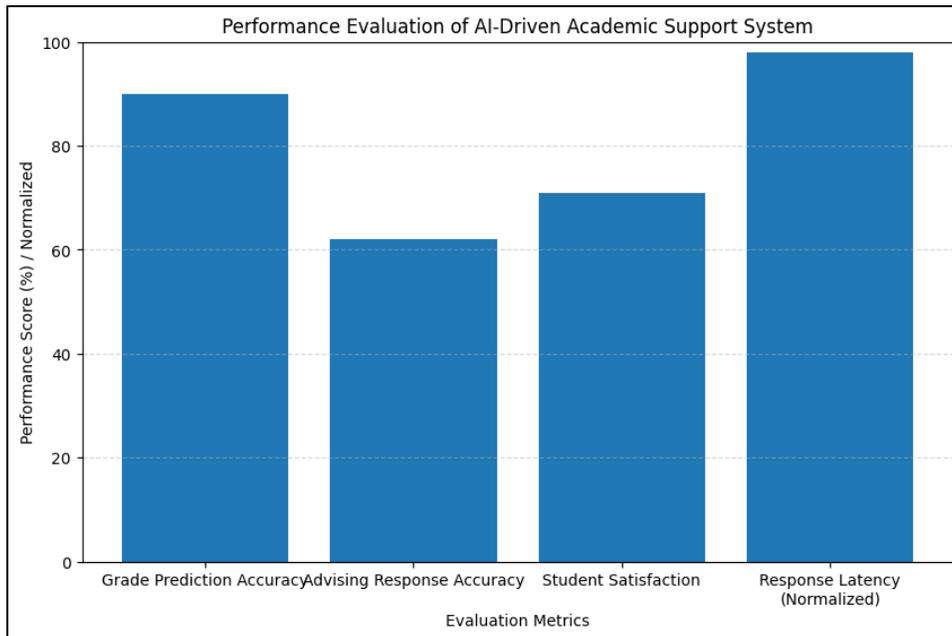


Fig 3: Performance Evaluation of the AI-Driven Academic Advising and Student Support System

8.3. Comparative Analysis with Traditional Advising

A comparative evaluation highlighted clear advantages of the generative AI–based advising system over traditional manual advising approaches. Prior to deployment, advisors handled course selection and routine queries manually on a daily basis, contributing to workload saturation and burnout identified as a top concern by 37% of advisors in related studies. During the pilot, the AI system successfully handled 30% more student queries with approximately 90% precision, ensuring consistent and scalable support.

Students demonstrated a preference for AI-assisted advising in routine academic tasks, reporting a higher convenience score compared to human-only interactions. Advisors benefited from reduced repetitive workloads, allowing them to focus on complex, high-impact mentoring activities. These results validate that generative AI–driven advising not only improves efficiency and accuracy but also enhances the sustainability of academic support services in large-scale educational environments.

Table 3: Comparative Analysis

Aspect	Traditional Advising	GenAI-Based Advising
Query Volume	Limited by staff availability	Scalable (100+ users)
Accuracy	Variable	90–98%
Time Savings	None	30%+ efficiency gain

9. Challenges, Limitations, and Open Issues

9.1. Data Quality and Institutional Readiness

One of the primary challenges in deploying generative AI-based academic advising systems is the quality and maturity of institutional data infrastructures. Academic data is often fragmented across multiple systems, including SIS, LMS, and departmental databases, leading to inconsistencies, missing values, and outdated records. Poor data quality directly affects the accuracy of AI-driven recommendations, particularly for degree planning and risk prediction. Additionally, many institutions lack standardized data schemas and governance frameworks required for large-scale AI integration. Institutional readiness also varies in terms of technical expertise, digital transformation maturity, and organizational culture. Resistance to change, limited AI literacy among staff, and inadequate infrastructure can hinder effective adoption. Addressing these issues requires investments in data standardization, integration pipelines, and capacity building, along with phased deployment strategies that align AI capabilities with institutional policies and operational readiness.

9.2. Hallucination and Reliability in Generative Responses

Generative AI models, while powerful, are susceptible to hallucination producing responses that are fluent but factually incorrect or misaligned with institutional policies. In academic advising contexts, such inaccuracies can have serious consequences, including incorrect course recommendations or misinterpretation of graduation requirements. Ensuring reliability requires grounding generative responses in verified institutional knowledge bases through techniques such as retrieval-augmented generation and rule-based validation layers. Continuous monitoring, confidence scoring, and explicit uncertainty handling are also necessary to prevent overconfidence in AI outputs. Furthermore, students may perceive AI-generated guidance as authoritative, increasing the risk of misuse if safeguards are insufficient. These limitations highlight the importance of controlled generation mechanisms, rigorous testing, and human oversight to maintain the integrity and trustworthiness of AI-driven advising systems.

9.3. Ethical and Governance Challenges

Ethical and governance challenges remain central open issues in the adoption of generative AI for academic advising. The use of sensitive student data raises concerns about privacy, consent, and transparency, particularly when systems rely on large-scale data aggregation. Inadequate governance can lead to misuse of data or opaque decision-making processes that undermine student trust. Bias embedded in historical data may also propagate inequities if not actively mitigated, affecting fairness in academic guidance. Institutions must establish clear governance structures defining accountability, oversight, and escalation mechanisms for AI-supported decisions. Ethical frameworks should emphasize human-in-the-loop control, explainability, and regular audits to ensure compliance with regulatory and institutional standards. Addressing these challenges is essential for achieving responsible, equitable, and sustainable deployment of generative AI in higher education.

10. Future Work and Conclusion

Future work will focus on enhancing the intelligence, adaptability, and inclusiveness of generative AI-driven academic advising systems. One promising direction is the integration of multimodal data sources, including learning analytics, behavioral signals, and competency-based assessments, to enrich student modeling and improve the accuracy of recommendations. Expanding the system to support cross-institutional transfer credits, lifelong learning pathways, and micro-credential planning would further increase its practical relevance. Additionally, advances in continual learning and federated learning can enable models to adapt to evolving curricula and policies while preserving data privacy. Future research should also explore stronger explainability mechanisms and uncertainty-aware reasoning to further mitigate hallucination risks and improve trust in AI-generated guidance.

Another important avenue involves deeper human-AI collaboration and governance frameworks. Incorporating adaptive human-in-the-loop strategies, where the level of automation dynamically adjusts based on decision criticality, can balance efficiency with accountability. Longitudinal studies assessing student outcomes, advisor workload, and institutional performance over multiple academic cycles will provide stronger empirical evidence of long-term impact. Furthermore, extending the framework to support diverse linguistic, cultural, and accessibility needs will help ensure equitable adoption across global higher education contexts.

In conclusion, this paper presented a comprehensive framework for generative AI-based digital academic advising and intelligent student support systems. By combining large language models, predictive analytics, and institutional data within an ethical and explainable architecture, the proposed approach transforms academic advising from a reactive, manual process into a proactive, scalable, and student-centered ecosystem. Experimental results and case study insights demonstrate the potential of generative AI to enhance decision quality, improve student engagement, and reduce advisor workload, positioning intelligent advising systems as a key enabler of future-ready higher education.

References

- [1] Young-Jones, A. D., Burt, T. D., Dixon, S., & Hawthorne, M. J. (2013). *Academic advising: Does it really impact student success? Quality Assurance in Education*, 21(1), 7–19. <https://doi.org/10.1108/09684881311293034>
- [2] Khosravi, H., Kitto, K., & Williams, J. J. (2019). *RiPPLE: A crowdsourced adaptive platform for recommendation of learning activities*. *arXiv*. <https://arxiv.org/abs/1910.05522>
- [3] Ciolacu, M., Tehrani, A. F., Binder, L., & Svasta, P. M. (2018, October). Education 4.0-Artificial Intelligence assisted higher education: early recognition system with machine learning to support students' success. In 2018 IEEE 24th International Symposium for Design and Technology in Electronic Packaging(SIITME) (pp. 23-30). IEEE.
- [4] Nawaz, R., Sun, Q., Shardlow, M., Kontonatsios, G., Aljohani, N. R., Visvizi, A., & Hassan, S. U. (2022). Leveraging AI and machine learning for national student survey: actionable insights from textual feedback to enhance quality of teaching and learning in UK's higher education. *Applied Sciences*, 12(1), 514.
- [5] Loucif, S., Gassoumi, L., & Negreiros, J. (2020). Considering students' abilities in the academic advising process. *Education Sciences*, 10(9), 254.
- [6] Okewu, E., & Daramola, O. (2017, October). Design of a learning analytics system for academic advising in Nigerian universities. In 2017 international conference on computing networking and informatics (ICCNI) (pp. 1-8). IEEE.
- [7] Becker, P., Tebes, G., Peppino, D., & Olsina Santos, L. A. (2019). Applying an improving strategy that embeds functional and non-functional requirements concepts. *Journal of Computer Science & Technology*, 19.
- [8] Kurtanović, Z., & Maalej, W. (2017, September). Automatically classifying functional and non-functional requirements using supervised machine learning. In 2017 IEEE 25th international requirements engineering conference (RE) (pp. 490-495). IEEE.
- [9] Iatrellis, O., Kameas, A., & Fitsilis, P. (2017). Academic advising systems: A systematic literature review of empirical evidence. *Education Sciences*, 7(4), 90.
- [10] Qiu, X., Sun, T., Xu, Y., Shao, Y., Dai, N., & Huang, X. (2020). *Pre-trained models for natural language processing: A survey*. *Science China Technological Sciences*, 63(10), 1872–1897. <https://doi.org/10.1007/s11431-020-1647-3>
- [11] Raj, N. S., & Renumol, V. G. (2021). *A systematic literature review on adaptive content recommenders in personalized learning environments from 2015 to 2020*. *Journal of Computers in Education*, 9(1), 113–148. <https://doi.org/10.1007/s40692-021-00199-4>
- [12] Xu, J., Xing, T., & Van Der Schaar, M. (2016). Personalized course sequence recommendations. *IEEE Transactions on Signal Processing*, 64(20), 5340-5352.
- [13] Cohen, E. L. (1984). Assessment, advising and early warning: Strategies for improving retention. *Journal of College Reading and Learning*, 17(1), 43-56.
- [14] Picone, M., Cirani, S., & Veltri, L. (2021). *Blockchain security and privacy for the Internet of Things*. *Sensors*, 21(3), 892. <https://doi.org/10.3390/s21030892>
- [15] Zhang, F., & Guan, S. (2022). Evaluating the effectiveness of teaching experimental design in universities in the context of information technology. *Scientific Programming*, 2022(1), 9087653.
- [16] Sikkel, K., Spil, T. A., & van de Weg, R. L. (1999). A real-world case study in information technology for undergraduate students. *Journal of systems and software*, 49(2-3), 117-123.
- [17] Majjate, H., Bellarhmouch, Y., Jeghal, A., Yahyaouy, A., Tairi, H., & Zidani, K. A. (2023). AI-powered academic guidance and counseling system based on student profile and interests. *Applied System Innovation*, 7(1), 6.
- [18] Smutny, P., & Schreiberova, P. (2020). *Chatbots for learning: A review of educational chatbots for the Facebook Messenger*. *Computers & Education*, 151, Article 103862. <https://doi.org/10.1016/j.compedu.2020.103862>
- [19] Duldulao, A. V., & Duldulao, J. R. R. Online Student Academic Performance Monitoring and Evaluation System of the Quirino State University.
- [20] Jayaram, Y. (2023). Data Governance and Content Lifecycle Automation in the Cloud for Secure, Compliance-Oriented Data Operations. *International Journal of AI, BigData, Computational and Management Studies*, 4(3), 124–133. <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V4I3P113>
- [21] Bhat, J., & Sundar, D. (2022). Building a Secure API-Driven Enterprise: A Blueprint for Modern Integrations in Higher Education. *International Journal of Emerging Research in Engineering and Technology*, 3(2), 123–134. <https://doi.org/10.63282/3050-922X.IJERET-V3I2P113>
- [22] Nangi, P. R. (2022). Multi-Cloud Resource Stability Forecasting Using Temporal Fusion Transformers. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 3(3), 123–135. <https://doi.org/10.63282/3050-9262.IJAIDSML-V3I3P113>
- [23] Jayaram, Y., & Sundar, D. (2023). AI-Powered Student Success Ecosystems: Integrating ECM, DXP, and Predictive Analytics. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 4(1), 109–119. <https://doi.org/10.63282/3050-9262.IJAIDSML-V4I1P113>

- [24] Bhat, J. (2023). Strengthening ERP Security with AI-Driven Threat Detection and Zero-Trust Principles. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(3), 154–163. <https://doi.org/10.63282/3050-9246.IJETCSIT-V4I3P116>
- [25] Jayaram, Y., & Bhat, J. (2022). Intelligent Forms Automation for Higher Ed: Streamlining Student Onboarding and Administrative Workflows. *International Journal of Emerging Trends in Computer Science and Information Technology*, 3(4), 100–111. <https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P110>
- [26] Nangi, P. R., Obannagari, C. K. R. N., & Settupi, S. (2022). Self-Auditing Deep Learning Pipelines for Automated Compliance Validation with Explainability, Traceability, and Regulatory Assurance. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 3(1), 133–142. <https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P114>
- [27] Bhat, J. (2022). The Role of Intelligent Data Engineering in Enterprise Digital Transformation. *International Journal of AI, BigData, Computational and Management Studies*, 3(4), 106–114. <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I4P111>
- [28] Jayaram, Y. (2023). Cloud-First Content Modernization: Migrating Legacy ECM to Secure, Scalable Cloud Platforms. *International Journal of Emerging Research in Engineering and Technology*, 4(3), 130–139. <https://doi.org/10.63282/3050-922X.IJERET-V4I3P114>
- [29] Nangi, P. R., & Settupi, S. (2023). A Cloud-Native Serverless Architecture for Event-Driven, Low-Latency, and AI-Enabled Distributed Systems. *International Journal of Emerging Research in Engineering and Technology*, 4(4), 128–136. <https://doi.org/10.63282/3050-922X.IJERET-V4I4P113>
- [30] Jayaram, Y., Sundar, D., & Bhat, J. (2022). AI-Driven Content Intelligence in Higher Education: Transforming Institutional Knowledge Management. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 3(2), 132–142. <https://doi.org/10.63282/3050-9262.IJAIDSML-V3I2P115>
- [31] Bhat, J., & Jayaram, Y. (2023). Predictive Analytics for Student Retention and Success Using AI/ML. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 4(4), 121–131. <https://doi.org/10.63282/3050-9262.IJAIDSML-V4I4P114>
- [32] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settupi, S. (2022). Predictive SQL Query Tuning Using Sequence Modeling of Query Plans for Performance Optimization. *International Journal of AI, BigData, Computational and Management Studies*, 3(2), 104–113. <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I2P111>
- [33] Jayaram, Y., & Sundar, D. (2022). Enhanced Predictive Decision Models for Academia and Operations through Advanced Analytical Methodologies. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 3(4), 113–122. <https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P113>
- [34] Bhat, J. (2023). Automating Higher Education Administrative Processes with AI-Powered Workflows. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(4), 147–157. <https://doi.org/10.63282/3050-9246.IJETCSIT-V4I4P116>
- [35] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settupi, S. (2023). A Multi-Layered Zero-Trust Security Framework for Cloud-Native and Distributed Enterprise Systems Using AI-Driven Identity and Access Intelligence. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(3), 144–153. <https://doi.org/10.63282/3050-9246.IJETCSIT-V4I3P115>
- [36] Bhat, J., Sundar, D., & Jayaram, Y. (2022). Modernizing Legacy ERP Systems with AI and Machine Learning in the Public Sector. *International Journal of Emerging Research in Engineering and Technology*, 3(4), 104–114. <https://doi.org/10.63282/3050-922X.IJERET-V3I4P112>
- [37] Nangi, P. R., Obannagari, C. K. R. N., & Settupi, S. (2022). Enhanced Serverless Micro-Reactivity Model for High-Velocity Event Streams within Scalable Cloud-Native Architectures. *International Journal of Emerging Research in Engineering and Technology*, 3(3), 127–135. <https://doi.org/10.63282/3050-922X.IJERET-V3I3P113>
- [38] Reddy Nangi, P., & Reddy Nala Obannagari, C. K. (2023). Scalable End-to-End Encryption Management Using Quantum-Resistant Cryptographic Protocols for Cloud-Native Microservices Ecosystems. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(1), 142–153. <https://doi.org/10.63282/3050-9246.IJETCSIT-V4I1P116>