IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AN AGENTIC AI FRAMEWORK FOR AUTOMATING ACADEMIC ADMINISTRATIVE TASKS IN HIGHER EDUCATION INSTITUTIONS

[1]Abhishek Kisan Pradhan
[1]Assistant Professor
[1]Department of Science & Technology
[1]GH Raisoni College of Engineering & Management, Nagpur, Maharashtra, India

Abstract: Higher education institutions (HEIs) handle a large volume of repetitive administrative tasks, such as responding to student queries, verifying academic records, scheduling, and document processing, that consume considerable staff time. Rule-based chatbots have partially automated such tasks, but they lack the multi-step reasoning and decision-making required for genuinely complex administrative workflows. This paper examines the application of Agentic Artificial Intelligence (Agentic AI), AI systems capable of planning, invoking external tools, taking autonomous actions, and verifying their own outputs, to the automation of academic administrative processes. We review the evolution of agentic architectures, from reasoning-and-acting paradigms to modern multi-agent orchestration frameworks, and propose a conceptual framework in which specialised agents handle distinct administrative functions (query resolution, records verification, scheduling, and escalation) under a coordinating orchestrator agent. We describe a proposed implementation approach using open-source agent-orchestration frameworks and synthetic academic data to avoid privacy risk, together with an evaluation methodology based on task-completion accuracy, response latency, and escalation correctness. The paper concludes with a discussion of feasibility, ethical considerations surrounding student data privacy, and directions for future empirical work. The proposed framework offers a low-cost, extensible pathway for resource-constrained HEIs, particularly in the Indian context, to reduce administrative overhead while retaining human oversight over sensitive decisions.

***Keywords* -** Agentic AI; Large Language Models; Multi-Agent Systems; Higher Education Administration; Academic Automation; Intelligent Agents

I. INTRODUCTION

Colleges and universities, particularly in India where student-to-administrative-staff ratios are often high, spend a substantial share of institutional effort on routine administrative work: answering repeated student queries about attendance and fee status, verifying eligibility for examinations, scheduling classes and re-examinations, and processing routine documentation. This burden diverts staff and faculty time away from teaching, mentoring, and research.

Early attempts to automate such work relied on rule-based or menu-driven chatbots that could answer narrowly scoped, single-turn questions but could not reason across multiple data sources or take a sequence of dependent actions. The emergence of Large Language Models (LLMs) improved natural-language understanding considerably, yet a standard LLM chatbot, prompted once and returning one answer, still cannot verify a claim against institutional records, decide when to escalate a borderline case to a human, or coordinate several sub-tasks toward a single administrative outcome.

Agentic AI refers to a class of LLM-based systems that go beyond single-turn response generation. An agentic system can decompose a goal into steps, decide which external tool or data source to consult at each step, take an action, observe the result, and revise its plan, often coordinating with other specialised agents before returning a final, self-checked response to the user. This shift, from answering to reasoning-and-acting, is what makes agentic AI a meaningfully different technology from earlier chatbot systems, and is the reason it is receiving significant research attention at present.

This paper explores how an agentic AI architecture could be applied specifically to academic administrative automation in higher education institutions. The specific objectives of this paper are to:

- Review the technical evolution of agentic AI, from reasoning-and-acting prompting to multi-agent orchestration frameworks, and the existing literature on chatbot-based automation in higher education administration.
- Propose a conceptual multi-agent framework in which an orchestrator agent delegates academic-administrative sub-tasks to specialised agents for query resolution, records verification, scheduling, and human escalation.
- Describe a low-cost, implementation-feasible approach suited to a typical Indian HEI setting, using open-source orchestration frameworks and synthetic data rather than real student records.
- Propose an evaluation methodology, and discuss feasibility, limitations, and the ethical and data-privacy considerations relevant to deploying such a system.

The remainder of this paper is organised as follows. Section 2 reviews related literature. Section 3 presents the proposed agentic framework and architecture. Section 4 describes a proposed implementation approach. Section 5 proposes an evaluation methodology. Section 6 discusses feasibility, limitations, and ethical considerations, and Section 7 concludes with directions for future work.

2. Literature Review

The literature relevant to this work spans two strands: (i) the technical evolution of agentic AI architectures, and (ii) the application of conversational AI to higher-education administration.

2.1 From Reasoning to Acting: Technical Foundations of Agentic AI

Yao et al. (2023) introduced ReAct, a prompting paradigm that interleaves reasoning traces with task-specific actions, allowing an LLM to update its plan, handle exceptions, and interface with external tools such as knowledge bases. ReAct is widely regarded as a foundational step toward agentic behaviour in LLMs, as it demonstrated that reasoning and acting are more effective when combined than when treated as separate capabilities.

Building on this foundation, Wang et al. (2024) present a comprehensive survey of LLM-based autonomous agents, proposing a unified conceptual framework consisting of a profiling module, a memory module, a planning module, and an action module. This framework is a useful reference point for designing domain-specific agents, including the administrative agents proposed later in this paper. Similarly, Xi et al. (2025) survey the rise of LLM-based agents and highlight open challenges around reliability, evaluation, and safe deployment, both of which are directly relevant to deploying such systems in a sensitive setting such as student administration.

Multi-agent orchestration has since become an active area of practical framework development. Wu et al. (2023) introduce AutoGen, an open-source framework in which customisable, conversable agents interact with each other, with humans, and with tools to accomplish tasks, and demonstrate its effectiveness across domains including mathematics, coding, and decision-making. Comparable open-source orchestration frameworks (e.g. LangChain and CrewAI) have since made multi-agent development accessible without requiring large research infrastructure, which is an important practical consideration for the implementation proposed in this paper.

Finally, since institution-specific administrative rules (attendance thresholds, fee-payment deadlines, examination eligibility criteria) are not part of a general-purpose LLM's training data, grounding agent responses in institutional documents is essential. Gao et al. (2023) survey Retrieval-Augmented Generation (RAG), a technique that allows an LLM to retrieve and condition its response on relevant external documents at inference time, reducing hallucination and keeping responses aligned to source policy documents.

2.2 Chatbots and AI in Higher Education Administration

A separate body of literature has examined chatbot adoption in higher education more broadly. Labadze et al. (2023), in a systematic literature review of 67 studies, find that AI chatbots benefit students primarily through homework and study assistance, personalised learning support, and skill development, while also noting their use by staff for administrative time-savings such as scheduling and routine student queries. However, this and related reviews consistently observe that most deployed systems are single-turn, intent-classification chatbots rather than agentic systems capable of multi-step reasoning, tool use, or autonomous cross-checking against institutional records.

This gap, between the technical capability of modern agentic AI architectures and their limited application to administrative automation specifically within higher education, motivates the framework proposed in this paper. To the authors' knowledge, there is limited published work that applies a multi-agent, tool-using architecture, as opposed to a single conversational chatbot, specifically to academic-administrative workflows in an Indian HEI context.

3. Agentic AI Framework for Academic Administration

This section proposes a conceptual multi-agent architecture in which distinct administrative responsibilities are assigned to specialised agents that operate under a coordinating orchestrator agent, rather than a single monolithic chatbot attempting every task. Figure 1 illustrates the proposed architecture.

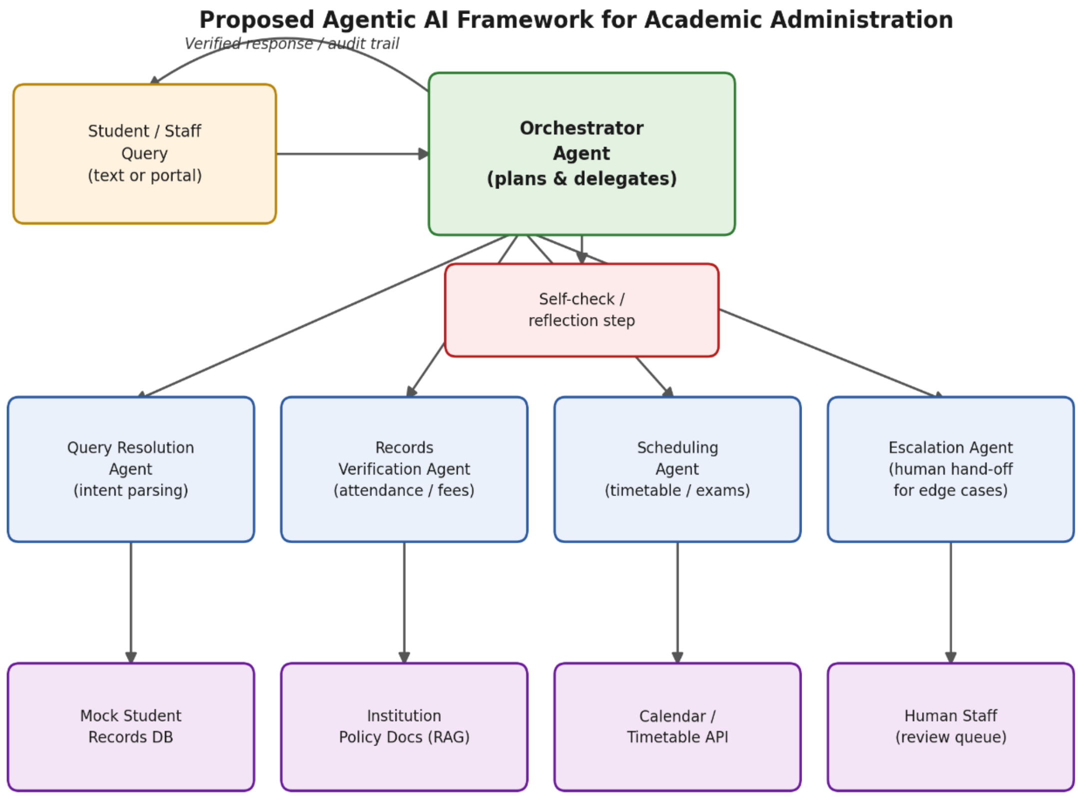

Figure 1. Proposed agentic AI architecture for academic administrative automation.

3.1 Components of the Architecture

Orchestrator Agent: receives the incoming student or staff query, decomposes it into sub-tasks, and delegates each sub-task to the relevant specialised agent. It also performs a final self-verification (reflection) step before returning a response, checking the drafted answer against the retrieved facts for consistency.

Query Resolution Agent: parses the natural-language query to identify intent (e.g., "Am I eligible to sit for the semester exam?") and the entities involved (student ID, course, semester).

Records Verification Agent: consults a records data source (attendance percentage, fee-payment status, backlog subjects) and institutional policy documents, retrieved via a RAG component, to check the relevant rule (e.g., minimum 75% attendance required for exam eligibility).

Scheduling Agent: interfaces with a calendar or timetable data source to answer or act on scheduling-related requests, such as re-examination dates or room allocation.

Escalation Agent: applies a confidence and risk threshold to decide whether a case is routine (safe to answer automatically) or borderline/sensitive (e.g., a disputed attendance record, a grievance, or any case involving personal data anomalies), in which case it routes the query to a human staff member along with the agent's reasoning trace, rather than allowing the system to decide autonomously.

This division of labour keeps each agent's prompt, tool access, and reasoning scope narrow and auditable, which is preferable to a single general-purpose agent both for reliability and for the institution's ability to review why a particular decision was made.

3.2 Illustrative Workflow

As an illustrative example: a student asks, "Am I eligible to appear for my semester-end examination?" The orchestrator agent identifies this as an eligibility query and delegates it to the Query Resolution Agent, which extracts the student's identifier and the relevant academic term. The Records Verification Agent then checks the student's attendance percentage and fee status against a data source, and retrieves the applicable eligibility rule from the institution's examination policy document via RAG. The orchestrator combines these results, drafts a response, and performs a self-check to confirm the drafted answer is consistent with the retrieved attendance figure and policy threshold. If the student's attendance is marginally below the cutoff, for instance, the Escalation Agent flags the case for manual review by the examination cell rather than issuing a final decision automatically, and the student instead receives an interim response indicating that the query has been forwarded for review. Every step is logged, giving the institution a complete audit trail of what data was consulted and what reasoning was applied.

This workflow demonstrates the central distinction from a conventional chatbot: the system does not merely retrieve an FAQ answer, it performs multi-step reasoning across two independent data sources, checks its own output, and knows when to defer to a human.

4. IMPLEMENTATION APPROACH

A key design goal of this proposal is feasibility within a typical Indian HEI setting, i.e., without dependence on large computing infrastructure or access to real, sensitive student data during development and evaluation. The following implementation approach is proposed:

- Orchestration framework: an open-source multi-agent framework such as AutoGen, LangChain, or CrewAI, all of which are free to use and can run on a standard laptop or a modest cloud instance.
- Underlying LLM: either a free-tier hosted LLM API for initial prototyping, or a small open-weight model (e.g., in the Llama or Mistral family) run locally, which avoids recurring API cost and keeps data on-premises, an important consideration for institutional data policy.
- Data: a synthetically generated dataset of mock student records (attendance percentages, fee status, backlog subjects) and de-identified institutional policy documents (attendance rules, examination eligibility criteria), constructed specifically so that no real student data is used during development or evaluation.
- Grounding: the institution's actual policy handbook (anonymised or with sensitive sections removed) can be indexed for the RAG component so that agent responses cite the correct, current rule rather than a generic or outdated one.
- Logging and audit trail: each agent action, tool call, and the final self-check decision should be logged, both to support the evaluation described in Section 5 and to give the institution a transparent record for any decision that is contested.

Suggested development stages are: (i) define the roles, prompts, and permitted tools of each agent; (ii) generate the synthetic dataset and policy corpus; (iii) implement the orchestration logic and inter-agent communication; (iv) integrate the RAG retrieval component; (v) implement the escalation threshold and logging; and (vi) test the end-to-end system on a set of representative synthetic queries before any pilot involving real users.

5. Evaluation Methodology

As this paper presents a conceptual framework and proposed implementation, this section describes a proposed evaluation methodology rather than reporting empirical results; the authors intend the framework to be evaluated once a working prototype has been built. Suggested metrics are as follows:

- Task-completion accuracy: the proportion of synthetic test queries for which the system's final answer matches the correct ground-truth answer, computed separately for each agent's task and for the end-to-end pipeline.
- Escalation precision and recall: whether the Escalation Agent correctly identifies borderline or sensitive cases that should be routed to a human, evaluated against a manually labelled set of edge cases.
- Response latency: end-to-end time from query submission to final response, since administrative usefulness depends on the system being no slower than, and ideally much faster than, the manual process.
- Consistency under self-check: the rate at which the orchestrator's reflection step catches and corrects an initially incorrect draft response before it reaches the user.
- Perceived usefulness: a small-scale structured survey of student and staff volunteers who interact with the prototype using synthetic scenarios, following the pattern used in prior chatbot-in-education studies (Labadze et al., 2023).

A useful baseline for comparison would be the time and accuracy of the existing manual process, and, where feasible, the performance of a simple single-turn LLM chatbot without agentic planning or tool use, to isolate the specific benefit of the multi-agent, tool-using design proposed here.

6. Discussion

6.1 Potential Advantages

An agentic administrative system of this kind offers round-the-clock availability, consistent application of institutional rules, and the potential to free administrative and faculty time for higher-value work, consistent with time-saving benefits reported for chatbot-based administrative tools in the broader higher-education literature (Labadze et al., 2023). Because each administrative decision is logged with its reasoning trace, the system could also improve the traceability and auditability of routine decisions compared to informal, undocumented manual handling.

6.2 Limitations

LLM-based agents remain susceptible to hallucination and reasoning errors, particularly in multi-step tasks, which is why the proposed design restricts each agent's scope and includes an explicit self-check and human-escalation path rather than allowing fully autonomous decisions on sensitive matters. The system's usefulness also depends on the quality and currency of the underlying institutional data and policy documents; if these are incomplete or outdated, the agents will reason from incorrect premises. Initial setup, defining agent roles, preparing the RAG corpus, and building the synthetic evaluation dataset, requires meaningful one-time effort before any benefit is realised.

6.3 Ethical and Data Privacy Considerations

Any real deployment of this framework would process personally identifiable student data (attendance, fees, academic records) and must therefore comply with applicable data-protection requirements, including India's Digital Personal Data Protection Act, and with the institution's own data-governance policy. For this reason, the development and evaluation approach proposed in Section 4 deliberately relies on synthetic data rather than real student records. A production deployment would additionally require strict access control, encryption of stored records, a clear data-retention policy, and, given the escalation design in Section 3, a guarantee that automated agents never make final decisions on grievances, disciplinary matters, or disputes without human sign-off.

6.4 Feasibility for Indian Higher Education Institutions

The proposed reliance on open-source orchestration frameworks and, optionally, locally hosted open-weight models makes this framework comparatively low-cost to prototype, which is relevant for resource-constrained Indian colleges. Support for regional-language queries (e.g., Marathi or Hindi) would be a natural and valuable extension, given that many students may be more comfortable querying administrative systems in their first language.

7. Conclusion and Future Work

This paper proposed a conceptual, multi-agent Agentic AI framework for automating routine academic-administrative tasks in higher education institutions, in which an orchestrator agent delegates to specialised agents for query resolution, records verification, scheduling, and human escalation, with a self-verification step and full audit logging. Unlike conventional single-turn chatbots, the proposed system is designed to reason across multiple data sources, take a sequence of dependent actions, and recognise when a decision should be deferred to a human, which addresses a gap identified in the existing higher-education chatbot literature.

The immediate direction for future work is to implement the proposed prototype using the tools and synthetic-data approach described in Section 4, and to carry out the evaluation described in Section 5 to obtain empirical results, which were intentionally not fabricated in this conceptual paper. Beyond this, future work could extend the framework to additional administrative functions (hostel and library administration, grievance triage), conduct a small institutional pilot study with appropriate ethics approval and student consent, benchmark the agentic design against a conventional single-turn chatbot baseline, and explore fine-tuning a small open-weight model on institution-specific data to reduce dependence on external APIs and improve support for regional languages.

References

1. Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997.
2. Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: Systematic literature review. International Journal of Educational Technology in Higher Education, 20, Article 56. https://doi.org/10.1186/s41239-023-00426-1
3. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J.-R. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18(6), Article 186345.

4. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Zhang, S., Zhu, E., Li, B., Jiang, L., Zhang, X., & Wang, C. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation framework. arXiv preprint arXiv:2308.08155.
5. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., et al. (2025). The rise and potential of large language model based agents: A survey. Science China Information Sciences, 68(2), Article 121101.
6. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. Proceedings of the 11th International Conference on Learning Representations (ICLR 2023).