ARTIFICIAL INTELLIGENCE, MULTI-AGENT SYSTEMS

Multi-Agent College Admission Process System: A Review and Proposed Framework

Revati Shinde[1], Kalyani Yadav[1], Aditi Vadd[1], Payal More[1], A. D. Patil[2]

[1]Student, Department of Computer Science and Engineering, Sanjay Ghodawat Institute, Kolhapur, Maharashtra, India
[2]Guide, Department of Computer Science and Engineering, Sanjay Ghodawat Institute, Kolhapur, Maharashtra, India

Published November 2026

Abstract

Admission offices in Indian colleges handle a large volume of repetitive enquiries about eligibility, fees, documents and Centralized Admission Process (CAP) rounds, and conversational artificial intelligence (AI) is increasingly proposed to absorb this load. This paper reviews fifteen works covering educational chatbots, university and admission assistants, retrieval-augmented generation (RAG), hallucination and abstention in large language models (LLMs), speech recognition, human-AI interaction guidelines and data-protection law. The review finds that educational chatbots are well studied for learning and FAQ support, and that retrieval reduces hallucination in generated answers. However, admission-facing systems rarely control which knowledge is allowed into answers, rarely decline when a candidate-specific record is missing, seldom model consent for personalised outreach, and do not distinguish drafted or simulated actions from real communication. Based on these gaps, the paper proposes a human-in-the-loop framework with four mechanisms: an approval-gated knowledge lifecycle, verified-record grounding with explicit abstention for CAP queries, a consent gate before any personalised message, and an explicit action-state model that labels drafts, previews and blocked actions. The framework is presented as a design proposal together with evaluation criteria for grounding, abstention, consent enforcement and record completeness, to guide future pilot studies in admission offices.

Keywords: abstention; admission chatbot; consent management; hallucination; human-in-the-loop AI; retrieval-augmented generation

1 Introduction

Admission season concentrates a large number of enquiries into a few weeks. In Maharashtra, admission to many technical programmes is conducted through Centralized Admission Process (CAP) rounds, and students and parents repeatedly ask about eligibility, fees, cut-offs, document deadlines and allotment status. Admission staff usually answer these questions by telephone, keep notes in spreadsheets and track callbacks in personal lists, so answers become inconsistent and follow-ups are missed.

Chatbots have been studied widely in education, for learning support, student assistance and administration [1–3]. In admissions, an AI virtual assistant that answered admitted students' questions by text message was associated with higher timely enrolment in a randomised study [4]. These results suggest that automated assistance can reduce staff load and help applicants.

Using large language models (LLMs) in admissions, however, brings three specific risks. First, LLMs can produce fluent but unsupported statements, a problem known as hallucination [5]. Second, candidate-specific questions, such as whether a seat was allotted in a given round, must be answered only from verified records; when the record is missing, the correct behaviour is to abstain [6]. Third,

an assistant that drafts messages can be mistaken for one that sends them, which conflicts with guidance that AI systems should make clear what they can do [7] and with legal requirements for consent when processing personal data [8].

This paper reviews the literature relevant to these risks and makes three contributions:

1. A thematic review of fifteen works on educational and admission chatbots, retrieval-augmented generation, hallucination and abstention, speech recognition, human-AI interaction and data protection.
2. A set of six research gaps that limit safe use of AI assistants in college admissions.
3. A proposed human-in-the-loop framework that addresses each gap, with evaluation criteria for future pilot studies.

Section 2 describes the review method. Literature review is presented in Section 3 and the research gaps are identified in Section 4. Section 5 describes the proposed solution. Section 6 discusses evaluation and challenges. Section 7 concludes.

2 Review methodology

A narrative literature review was conducted. Sources were searched in IEEE Xplore, the ACM Digital Library, the ACL Anthology, SpringerLink, Elsevier ScienceDirect, Frontiers and the Proceedings of Machine Learning Research, together with Indian legislation. Search terms included "educational chatbot", "university chatbot", "admission assistant", "retrieval-augmented generation", "hallucination", "abstention", "speech recognition", "human-AI interaction" and "data protection consent".

Works were included if they were peer-reviewed or widely cited preprints published between 2017 and 2025, and if they addressed at least one of four review questions: RQ1, what tasks educational and admission chatbots perform; RQ2, how answers are grounded in reliable information; RQ3, how missing or uncertain information is handled; and RQ4, how consent and accountability for actions are handled. Fifteen works were selected and grouped into six themes, summarised in Table 1.

3 Literature review

3.1 Chatbots in education

Wollny et al. [1] conducted a systematic literature review of chatbots in education and classified their objectives and pedagogical roles, such as supporting learning, assisting students and mentoring. They concluded that the field is still at an early stage and that evidence on effectiveness is limited. Okonkwo and Ade-Ibijola [2] reviewed the applications of chatbots in education and pointed out the advantages such as quick access to information, support for multiple users simultaneously, and challenges such as ethics, evaluation and integration with institutional data .

Kuhail et al. [3] reviewed how students interact with educational chatbots and found that most systems were text-based and followed predefined conversation flows, with relatively few studies of speech-based interaction.

3.2 University and admission assistants

Page and Gehlbach [4] evaluated an AI virtual assistant at Georgia State University that answered admitted students' questions by text message, and reported improved timely enrolment in a

randomised trial. Ranoliya et al. [9] built a chatbot for university FAQs using Artificial Intelligence Markup Language (AIML) pattern matching combined with semantic analysis. Hien et al. [10] presented FIT-EBot, which used intent classification and a knowledge base to answer administrative and learning questions in a university faculty. These systems confirm demand for automated support, but their answers depend on authored patterns or curated content, and none verifies candidate-specific records before answering.

3.3 Retrieval-augmented generation

Lewis et al. [11] introduced retrieval-augmented generation (RAG), which conditions a generator on passages retrieved from an external index and improves performance on knowledge-intensive tasks while allowing knowledge to be updated without retraining. Karpukhin et al. [12] showed that dense passage retrieval with a dual encoder can outperform sparse keyword retrieval for open-domain question answering. Gao et al. [13] surveyed the development of RAG from naive to advanced and modular designs and listed open challenges, including robustness and evaluation. Shuster et al. [14] found that retrieval augmentation reduces hallucination in knowledge-grounded conversation. A common assumption across these works is that the retrieval corpus is trustworthy; none defines how content is reviewed, approved or withdrawn before it can be retrieved.

3.4 Hallucination and abstention

Ji et al. [5] surveyed hallucination in natural language generation, distinguishing intrinsic hallucination, which contradicts the source, from extrinsic hallucination, which cannot be verified from it, and reviewed mitigation methods. Wen et al. [6] surveyed abstention in LLMs, framing it from the perspectives of the query, the model and human values, and argued that refusing to answer is often the safer choice. Most abstention methods rely on model confidence or uncertainty estimates, whereas admission queries often require abstention based simply on whether a verified record exists.

3.5 Speech interfaces

Radford et al. [15] trained a speech recognition model on large-scale weakly supervised data and reported robust zero-shot performance across many languages and conditions. Recognition errors nevertheless remain, especially in noisy audio and lower-resource languages, and in a voice assistant they can propagate to summaries and outcome labels constructed from transcripts.

3.6 Human-AI interaction and data privacy

Amershi et al. [7] proposed eighteen guidelines for human-AI interaction, e.g. making clear what the system can do and how good it can do, and enabling efficient correction and dismissal. India The Digital Personal Data Protection Act, 2023 [8] recognises consent as a basis for processing digital personal data and grants rights to individuals in relation to their data. Neither source describes technical mechanisms for an admission assistant, though together they define what such a system is expected to do.

3.7 Comparative analysis

Table 1 summarises the reviewed works against the four review questions. On RQ1, the education reviews [1–3] and the university systems [4, 9, 10] show that chatbots are used mainly to answer frequent questions, guide students through procedures and support learning. Admission-specific

evidence is limited to a small number of deployments, of which the randomised study in [4] is the strongest.

On RQ2, the older university chatbots ground answers in hand-written patterns or intent-response pairs [9, 10]. This keeps answers controlled but makes them rigid and costly to update. Retrieval-based approaches [11–14] are more flexible and reduce unsupported content, but they move the responsibility for correctness to the quality of the indexed corpus, which none of the reviewed works governs.

On RQ3, abstention is studied as a model behaviour driven by uncertainty [6], and hallucination mitigation is likewise model-centred [5]. None of the education or admission systems describes how the assistant should respond when a required record is absent. On RQ4, accountability and consent appear only in general guidance [7] and in law [8]; no reviewed system turns them into concrete checks inside the assistant.

The literature supplies the parts (retrieval, abstention concepts, interaction guidelines and legal consent requirements), but no reviewed work combines them for the workflow of an admission office. Section 4 sets out the gaps that follow from this.

Table 1. Summary of reviewed literature.

Ref.	Focus and approach	Limitation for admissions use
[1]	Systematic review of chatbots in education; roles and objectives	Learning-centred; little on administrative records or consent
[2]	Review of chatbot applications, benefits and challenges in education	Notes data-integration and ethics issues but proposes no controls
[3]	Review of interaction design of educational chatbots	Mostly text, flow-based bots; voice use rarely studied
[4]	AI text-message assistant for admitted students; randomised trial	Curated answers; no record-level verification or abstention
[5]	Survey of hallucination in natural language generation	Mitigation is model-centred, not record-centred
[6]	Survey of abstention in large language models	Abstention based on model uncertainty, not data presence
[7]	Eighteen guidelines for human-AI interaction	General guidance; not applied to admission workflows
[8]	Digital Personal Data Protection Act, 2023	Legal duty of consent; no technical mechanism specified
[9]	AIML pattern matching with semantic analysis for university FAQs	Static rules; answers limited to authored patterns
[10]	Intent-based chatbot for administrative and learning support	No knowledge approval workflow or candidate data
[11]	Retrieval-augmented generation over an external document index	Assumes the indexed corpus is trustworthy
[12]	Dense passage retrieval with a dual encoder	Retrieval quality only; no governance of sources
[13]	Survey of naive, advanced and modular RAG	Governance of sources treated as open challenge
[14]	Retrieval augmentation in knowledge-grounded dialogue	Reduces but does not remove hallucination

Ref.	Focus and approach	Limitation for admissions use
[15]	Large-scale weakly supervised speech recognition	Recognition errors persist under noise and in some languages

4 Research gap

Measured against the four review questions, the reviewed works leave six gaps that limit the safe use of AI assistants in college admissions.

G1 Ungoverned knowledge. RAG improves factual grounding [11, 14], but the reviewed works assume a trustworthy corpus [13]. Admission information such as fees and deadlines changes every academic year, yet no reviewed system provides an approval and withdrawal lifecycle for the knowledge an assistant may use.

G2 No record-level abstention. Abstention research is model-centred and confidence-based [6]. Candidate-specific CAP questions instead need deterministic abstention: if no verified record exists, the assistant must say so. University chatbots [9, 10] do not address candidate records at all.

G3 Consent not modelled. Educational chatbot reviews focus on pedagogy and interaction [1–3]. Personalised outreach to applicants requires consent [8], but no reviewed system checks consent before generating a candidate-specific message.

G4 Real and simulated actions blurred. Guidelines ask AI systems to make clear what they can do [7], yet the reviewed systems do not distinguish a drafted or previewed message from one actually sent, or a blocked campaign from a completed one.

G5 Speech errors propagate. Voice interaction is rarely studied in educational chatbots [3], and recognition errors [15] can silently corrupt summaries and follow-up decisions.

G6 Weak follow-up continuity. Existing assistants answer questions [4, 9, 10] but do not produce structured outcomes, such as dispositions and follow-up tasks, that counsellors can act on.

A further cross-cutting gap is evaluation: the reviewed works measure answer quality or user satisfaction, but not governance properties such as whether unapproved content leaks into answers or whether consent is enforced.

5 Proposed solution

5.1 Framework overview

We propose a human-in-the-loop framework in which the language model generates conversation, but governed data decides what is true and what has happened. As shown in Figure 1, admissions staff and applicants interact through a governance and service layer that handles authentication, the consent gate, action-state labels and audit logging. Three assistants sit behind this layer: an admission assistant for general questions, a CAP assistant for candidate-specific questions, and a messaging assistant for personalised replies. All three read from and write to a governed data layer consisting of a record store, which holds status, consent, verified CAP facts and audit events, and a vector index, which holds embeddings of approved knowledge only.

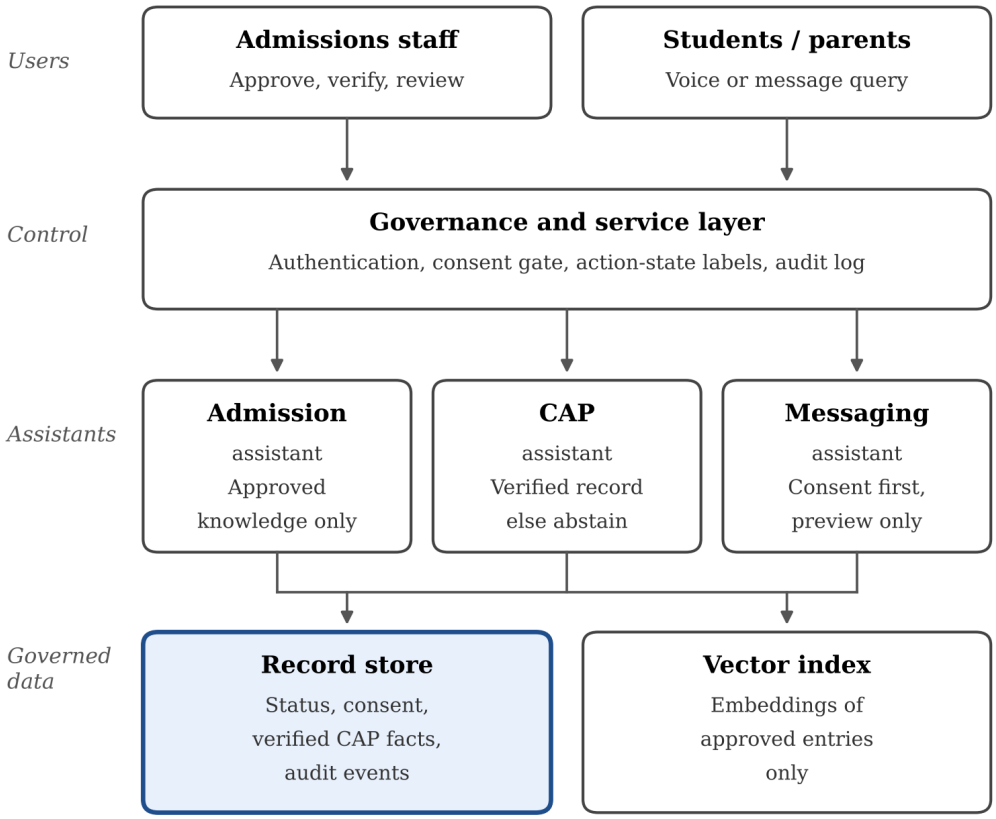


Figure 1. Proposed grounded, consent-aware admission assistant framework.

5.2 Approval-gated knowledge lifecycle (G1)

Each knowledge entry, such as a fee rule or a document deadline, carries a category, language, academic year and official source. New entries are saved as drafts and cannot be used in answers. Only after an administrator reviews and approves an entry is it embedded into the vector index. If indexing fails, the entry remains a draft and the failure is reported. At answer time, retrieved entries are checked again against the record store, so an entry withdrawn by staff stops being used immediately.

5.3 Verified-record grounding with abstention (G2)

The CAP assistant answers candidate-specific questions only from facts that administrators have recorded with an official source and marked as verified. The question is first mapped to the fact it requests, such as the allotment result for a given round. If a verified fact exists, the assistant states it with its source. If the question is general, such as round dates, the assistant uses approved knowledge. Otherwise it replies that the information is not yet recorded. The language model phrases the reply but never supplies the fact. Before use by applicants directly, candidate selection must be replaced by verified identity resolution.

5.4 Consent-gated, preview-only messaging (G3)

Before generating any candidate-specific message, the messaging assistant checks for a granted consent record. Without it, the request returns a consent-required result and nothing is generated. With consent, the reply is grounded in approved knowledge, filled into an English, Hindi or Marathi

template and saved as a preview for staff review. In the first phase no send or schedule path exists, so the system cannot claim delivery.

5.5 Explicit action-state model (G4)

Every object that represents an action carries a visible, stored state. Knowledge moves from draft to approved. Calling lists move from imported draft to authorised to campaign, whose start is shown as blocked until a telephony provider is integrated. Messages remain previews. Only events that really happened, such as a completed voice session or a created task, are marked completed. Staff therefore cannot mistake a simulated workflow for real communication. Figure 2 shows the four lifecycles; in each, the highlighted or dashed final state makes the limits of the current phase visible to staff.

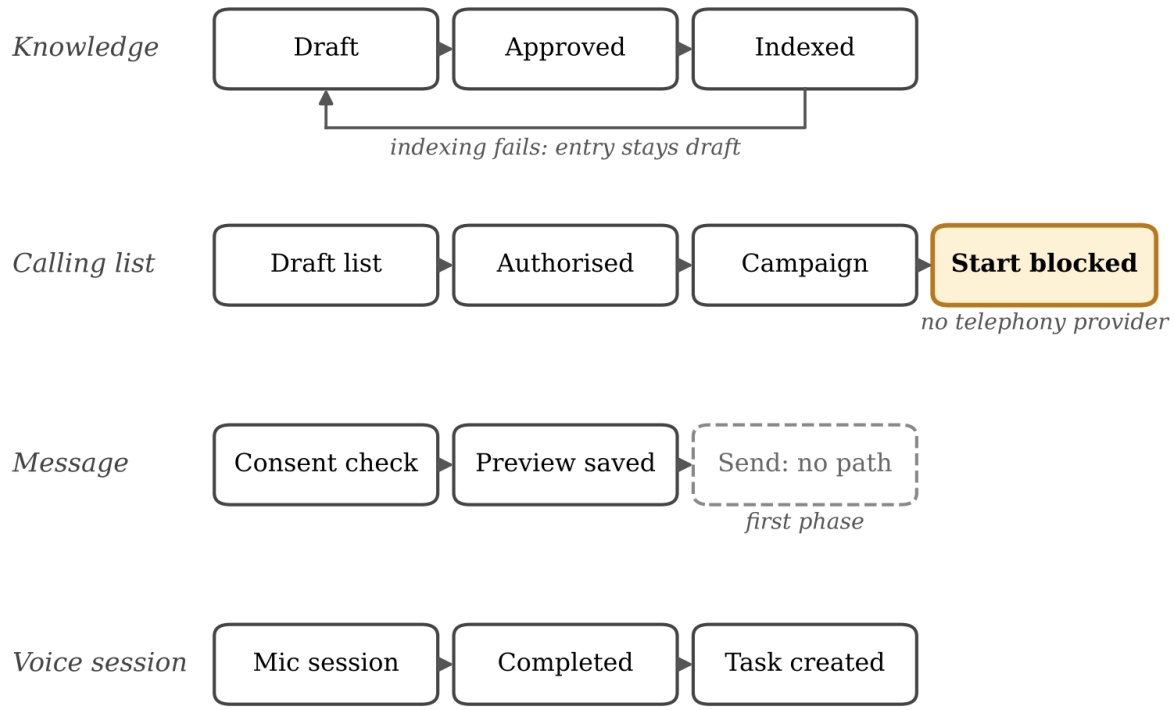


Figure 2. Action-state lifecycles for knowledge, calling lists, messages and voice sessions.

5.6 Structured records and follow-up (G5, G6)

When a voice session ends, its ordered transcript is stored with a summary, caller type, disposition and next action. Dispositions are normalised to Interested, Not Interested, Follow-up, Callback, Provisional or Pending, and Follow-up or Callback outcomes automatically create a task for a counsellor. Transcripts are labelled as speech recognition output, and summaries remain open to staff review, so recognition errors can be caught before they drive decisions. Table 2 maps each gap to its mechanism.

5.7 Illustrative workflow

In the first of two scenarios, a parent asks the admission assistant about the documents required for direct second-year admission. The assistant retrieves matching entries, discards any that are not

approved for the current academic year, and answers from what remains. If no approved entry covers the question, it says so and the question is logged as unresolved, so staff can add an approved entry for future enquiries.

In the second, a candidate asks the CAP assistant whether a seat was allotted in the second round. If staff have recorded a verified allotment fact for that candidate and round, the assistant states it together with the official source. If not, the assistant replies that the result is not yet recorded, rather than guessing from general cut-off information. In both scenarios, the conversation ends with a stored transcript, a disposition and, where needed, a follow-up task for a counsellor.

Table 2. Mapping of research gaps to proposed mechanisms.

Research gap	Proposed mechanism
G1 Ungoverned knowledge	Approval-gated knowledge lifecycle; only approved entries indexed
G2 No record-level abstention	Verified-record grounding; answer 'not recorded' when fact absent
G3 Consent not modelled	Consent gate before any personalised message
G4 Real and simulated actions blurred	Explicit action states: draft, approved, preview, blocked, completed
G5 Speech errors propagate	Transcripts labelled as recognised speech; staff review of summaries
G6 Weak follow-up continuity	Structured disposition and automatic follow-up task per conversation

6 Discussion

The framework moves trust decisions out of the language model and into data that the model cannot change. Approval, verification, consent and action state become database facts, so the framework's safety properties can be tested directly rather than depending on prompt instructions that a model may not follow.

6.1 Evaluation criteria

To address the evaluation gap identified in Section 4, Table 3 lists criteria for a pilot study. Each criterion is tied to the gap it tests. The deterministic criteria, such as draft leakage, consent enforcement and state truthfulness, can be checked with automated test cases. Answer grounding and transcript accuracy require judgement and should be assessed by two admissions staff members against the source entries, with disagreements resolved by discussion. Baselines should be established in a first pilot before numeric targets are set.

Table 3. Proposed evaluation criteria.

Criterion	Definition	Gap
Grounded answer rate	Share of answers traceable to an approved knowledge entry	G1
Draft leakage	Answers that use draft or withdrawn content (target: zero)	G1
Correct abstention	CAP questions without a verified fact that receive 'not recorded'	G2
False abstention	CAP questions with a verified fact that are wrongly declined	G2
Consent enforcement	Share of previews correctly blocked when consent is not granted	G3
State truthfulness	Cases where the interface implies a call or message was sent (target: zero)	G4

Criterion	Definition	Gap
Transcript accuracy	Agreement of summaries and dispositions with staff review	G5
Record completeness	Sessions saved with transcript, summary and disposition	G6
Task creation	Follow-up and Callback outcomes that create a visible task	G6

For comparison, the admission assistant should also be tested in two baseline configurations: the same language model answering without retrieval, and retrieval over all entries without the approval filter. The difference between these baselines and the full framework isolates the contribution of approval gating from that of retrieval itself.

6.2 Expected benefits

For applicants, the framework aims to give consistent answers that reflect the current academic year and to avoid confident but wrong statements about individual results. For admission staff, it aims to cut repetitive calls and keep every conversation and follow-up in one place, with the source of each answer traceable. For the institution, it offers an auditable record of what information was given, when consent was obtained and which actions actually took place.

6.3 Challenges and open issues

Verified identity resolution is required before applicants can query their own results, for example by combining an application number with a one-time password. Speech recognition quality in Marathi and Hindi needs local evaluation, since recognition errors can affect summaries and dispositions [15]. Approved knowledge must be reviewed by staff every academic year, which adds an administrative task. Retention, deletion and access review for applicant data must be defined in line with [8]. Integration with telephony and messaging providers will require careful handling of failures and delivery confirmation. Finally, the framework is designed to support counsellors and does not make admission decisions.

7 Conclusion

This paper reviewed fifteen works on educational and admission chatbots, retrieval-augmented generation, hallucination and abstention, speech recognition, human-AI interaction and data protection. The review identified six gaps: ungoverned knowledge, lack of record-level abstention, missing consent modelling, blurred real and simulated actions, propagation of speech errors and weak follow-up continuity. To address them, we proposed a human-in-the-loop framework combining approval-gated knowledge, verified-record abstention, consent-gated preview messaging and an explicit action-state model.

Future work will evaluate the framework in an admission season using the proposed criteria, add verified identity resolution, integrate telephony and messaging providers with delivery confirmation, display source citations alongside answers, and extend the framework to campuses with several institutions, each maintaining its own approved knowledge.

Acknowledgment

The authors thank their guide, A. D. Patil, and the Department of Computer Science and Engineering, Sanjay Ghodawat Institute, Kolhapur, for guidance and support throughout this work.

References

[1] Wollny S, Schneider J, Di Mitri D, Weidlich J, Rittberger M, Drachsler H. Are we there yet? A systematic literature review on chatbots in education. *Front Artif Intell*. 2021;4:654924. https://doi.org/10.3389/frai.2021.654924

[2] Okonkwo CW, Ade-Ibijola A. Chatbots applications in education: A systematic review. *Comput Educ Artif Intell*. 2021;2:100033. https://doi.org/10.1016/j.caeai.2021.100033

[3] Kuhail MA, Alturki N, Alramlawi S, Alhejori K. Interacting with educational chatbots: A systematic review. *Educ Inf Technol*. 2023;28(1):973–1018.

[4] Page LC, Gehlbach H. How an artificially intelligent virtual assistant helps students navigate the road to college. *AERA Open*. 2017;3(4). https://doi.org/10.1177/2332858417749220

[5] Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of hallucination in natural language generation. *ACM Comput Surv*. 2023;55(12):248. https://doi.org/10.1145/3571730

[6] Wen B, Yao J, Feng S, Xu C, Tsvetkov Y, Howe B, et al. Know your limits: A survey of abstention in large language models. *Trans Assoc Comput Linguist*. 2025;13:529–556. https://doi.org/10.1162/tacl_a_00754

[7] Amershi S, Weld D, Vorvoreanu M, Fourney A, Nushi B, Collisson P, et al. Guidelines for human-AI interaction. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*; 2019; Glasgow, UK. p. 1–13. https://doi.org/10.1145/3290605.3300233

[8] Government of India. *The Digital Personal Data Protection Act, 2023 (Act No*. 22 of 2023). New Delhi: Ministry of Law and Justice; 2023 Aug.

[9] Ranoliya BR, Raghuwanshi N, Singh S. Chatbot for university related FAQs. In: *Proceedings of the International Conference on Advances in Computing, Communications and Informatics (ICACCI)*; 2017; Udupi, India. p. 1525–1530. https://doi.org/10.1109/ICACCI.2017.8126057

[10] Hien HT, Cuong PN, Nam LNH, Nhung HLTK, Thang LD. Intelligent assistants in higher-education environments: The FIT-EBot, a chatbot for administrative and learning support. In: *Proceedings of the 9th International Symposium on Information and Communication Technology (SoICT)*; 2018; Danang, Vietnam. p. 69–76. https://doi.org/10.1145/3287921.3287937

[11] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In: *Advances in Neural Information Processing Systems (NeurIPS)*; 2020. Vol. 33, p. 9459–9474.

[12] Karpukhin V, Oguz B, Min S, Lewis P, Wu L, Edunov S, et al. Dense passage retrieval for open-domain question answering. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*; 2020. p. 6769–6781. https://doi.org/10.18653/v1/2020.emnlp-main.550

[13] Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, et al. Retrieval-augmented generation for large language models: A survey [Preprint]. *arXiv*:2312.10997; 2023.

[14] Shuster K, Poff S, Chen M, Kiela D, Weston J. Retrieval augmentation reduces hallucination in conversation. In: *Findings of the Association for Computational Linguistics: EMNLP 2021*; 2021. p. 3784–3803. https://doi.org/10.18653/v1/2021.findings-emnlp.320

[15] Radford A, Kim JW, Xu T, Brockman G, McLeavey C, Sutskever I. Robust speech recognition via large-scale weak supervision. In: *Proceedings of the 40th International Conference on Machine Learning (ICML)*; 2023. PMLR 202, p. 28492–28518.