Vol-3 Issue-1 2025
Scientific Research Journal of Science, Engineering and Technology
ISSN: 2584-0584, Peer Reviewed Journal

AI-Powered Meeting Assistant: An LLM-Centric, Agentic AI Approach for Automating Post-Meeting Workflows

Nikit Patel, Research Scholar, Computer Science, GLS University, Ahmedabad, India.
Kaushal Patel, Software Engineering, Electronic Communications, System Level Solutions: SLS, Anand, India

Abstract

Meetings are critical for collaborative decision-making, yet their outcomes are often underutilized due to inefficiencies in capturing, organizing, and tracking discussions. Traditional approaches to meeting documentation—manual note-taking or transcription-based solutions—fail to provide actionable insights, frequently leading to lost information, missed deadlines, and lack of accountability. With the advent of Large Language Models (LLMs) and agentic artificial intelligence (AI), it is now possible to design systems that not only document but also act upon meeting outcomes. This paper presents an LLM-centric AI-powered meeting assistant architecture that automates the complete post-meeting workflow. The system ingests audio or text transcripts, preprocesses them for clarity, and leverages LLMs to generate structured meeting summaries, extract actionable tasks, and create formal Minutes of Meeting (MoM). These outputs are converted into JSON format for seamless integration with task trackers, notification systems, and centralized dashboards. By adopting an agentic AI approach, the system enables proactive follow-ups, real-time reporting, and task completion monitoring through both tracker APIs and email-based confirmations. We provide a detailed literature review of speech-to-text technologies, LLM-driven meeting automation, and workflow orchestration, followed by a comprehensive description of the system's architecture and methodology. The benefits, limitations, and challenges—including speaker diarization, overlapping speech, and task misclassification—are critically examined. The study highlights the transformative potential of agentic AI for enterprise productivity, while emphasizing ethical considerations and the importance of human oversight.

Keywords

Meeting Automation; Large Language Models; Agentic AI; Speech-to-Text; Action Item Extraction; Task Tracking; Minutes of Meeting; Workflow Orchestration; Enterprise Productivity

---

1. Introduction

Meetings remain an indispensable element of organizational life. From strategic planning to project updates, collaborative discussions shape decision-making, allocate responsibilities, and set organizational direction. However, the productivity of meetings is frequently undermined by poor documentation, ambiguous follow-up actions, and ineffective tracking of outcomes (Rennard et al., 2023). Employees spend significant time revisiting meeting transcripts or email threads, while managers struggle with accountability gaps in task execution. Traditional solutions—manual note-taking, standalone transcription tools, or meeting minutes prepared by participants—rarely provide the structured, actionable intelligence required for effective follow-through.

Recent advances in Large Language Models (LLMs) have created new opportunities to transform meeting workflows (Laskar et al., 2023). LLMs excel at text comprehension, summarization, and structured generation, making them suitable for extracting insights from unstructured meeting transcripts (Golia & Kalita, 2023; Jin et al., 2025). Beyond this, the rise of agentic AI—systems capable of reasoning, acting, and integrating across APIs—enables the automation of downstream workflows, including task assignment, follow-up communication, and performance reporting (Asthana et al., 2025).

This research proposes and evaluates an LLM-centric meeting assistant architecture designed to automate the entire post-meeting lifecycle. By leveraging speech-to-text systems for input, LLMs for text understanding and generation, and APIs for task tracking and notifications, the system serves as an intelligent mediator between meeting discussions and organizational execution (Tan et al., 2023). The approach is not merely reactive but agentic, meaning the assistant proactively engages with data, monitors task completion, and ensures stakeholders remain aligned.

The paper proceeds as follows: Section 2 presents a literature review of existing research and technologies relevant to speech-to-text, meeting summarization, and agentic AI. Section 3 introduces the system architecture and provides a descriptive overview of the accompanying diagram. Section 4 outlines the methodology in detail. Section 5 discusses benefits, while Section 6 analyzes challenges and limitations. Section 7 explores future directions for research and deployment, followed by conclusions in Section 8.

2. Literature Review

2.1 Speech-to-Text and Automatic Speech Recognition

Speech-to-text (STT) technologies provide the foundational layer for meeting assistants by transforming audio recordings into textual transcripts. Early systems relied on Hidden Markov Models (HMMs) and n-gram language models, which often struggled with variability in accents, noise, and domain-specific terminology. Modern systems employ deep neural networks and self-supervised learning, as seen in Whisper (OpenAI), which offers multilingual support and speaker-aware transcription. Similarly, Google Speech-to-Text and AWS Transcribe are enterprise-grade services that emphasize scalability and integration. Despite progress, challenges remain in speaker diarization (distinguishing between multiple speakers), handling overlapping speech, and addressing language barriers (Schneider et al., 2025).

2.2 Meeting Summarization

Summarization has been studied extensively, evolving from extractive methods—where key sentences are selected—to abstractive approaches enabled by transformer-based models such as BERT and GPT (Feng et al., 2020; Zhu et al., 2020). Research in this domain emphasizes generating concise, coherent summaries that preserve decisions and critical points while excluding redundancies. Studies highlight that effective summarization must go beyond simple transcript compression to capture decisions, questions, and task allocations (Jin et al., 2025; Rennard et al., 2023).

2.3 Action Item Extraction

Task extraction from meeting dialogues requires identifying actionable statements, assigning responsible individuals, and inferring deadlines. Techniques range from rule-based approaches using linguistic patterns (e.g., imperative verbs such as “prepare,” “submit,” or “finalize”) to advanced neural architectures trained on annotated corpora. Recent work demonstrates that LLMs with prompt engineering can outperform traditional classifiers by directly generating structured task representations (Golia & Kalita, 2023). However, challenges such as vague instructions (“let’s work on this later”) or ambiguous ownership persist (Sadia et al., 2025).

2.4 LLMs in Enterprise Productivity

The rise of GPT-class models has spurred their application in enterprise productivity tools. Studies report successful use of LLMs in generating business reports, drafting emails, and summarizing technical documents (Laskar et al., 2023; Asthana et al., 2025). The unique advantage of LLMs lies in their few-shot and zero-shot learning ability, enabling them to generalize across tasks without extensive retraining. Yet, risks of hallucination, context truncation, and biases necessitate robust evaluation frameworks (Tan et al., 2023).

2.5 Agentic AI and Workflow Automation

Agentic AI refers to AI systems that are not limited to static predictions but can act autonomously, orchestrate multi-step workflows, and integrate with external tools. In meeting assistants, agentic AI enables proactive follow-ups (e.g., sending reminders before deadlines), dynamic adaptation (e.g., clarifying ambiguous instructions), and API-driven task management (Asthana et al., 2025). Literature highlights the growing interest in multi-agent systems and autonomous task orchestration, though most implementations remain experimental (Li et al., 2019).

2.6 Visualization and Reporting

Research in human-computer interaction emphasizes the role of dashboards in decision-making. Effective visualization supports awareness of task progress, overdue items, and overall meeting productivity. Integration with business intelligence platforms further enhances usability, though concerns of information overload and data privacy are frequently discussed (Deng et al., 2023).

This review highlights a convergence of fields—speech recognition, summarization, task extraction, and agentic AI—that collectively underpin the proposed meeting assistant architecture.

3. System Architecture

The proposed architecture integrates LLMs with a structured workflow to automate post-meeting activities (Laskar et al., 2023).

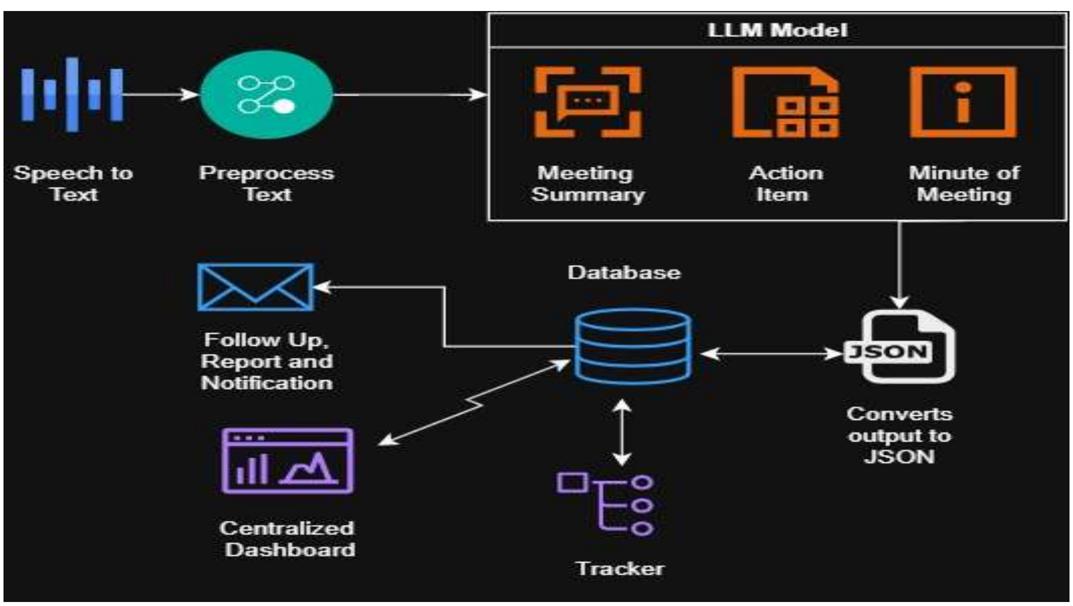

Figure 1. AI-powered meeting assistant workflow

Image Description:
The figure depicts the AI-powered meeting assistant workflow. Speech input is first converted into text, which undergoes preprocessing to remove noise, filler words, and inconsistencies. This processed transcript is then passed into the LLM model, which generates three key outputs: meeting summary, action items, and minutes of meeting (MoM). These outputs are subsequently converted into JSON format for structured integration. The JSON output is stored in a central database, which acts as the hub for downstream processes. The database connects to a task tracker (via APIs), a centralized dashboard (for visualization), and a notification system that emails stakeholders with summaries and assigned tasks. Completion tracking occurs either through updates from the tracker or through email confirmations parsed by the LLM. The entire architecture ensures seamless flow from raw meeting input to actionable organizational outcomes.

This modular architecture reflects a deliberate design choice: keeping the LLM as the central reasoning engine while ensuring interoperability with external enterprise systems (Asthana et al., 2025).

4. Methodology

The methodology involves several interlinked components designed to deliver a seamless end-to-end solution.

The process begins with speech-to-text conversion, where raw audio recordings are transcribed using ASR technologies such as Whisper or enterprise-grade solutions like AWS Transcribe (Schneider et al., 2025). Preprocessing then ensures the transcripts are usable: filler words are eliminated, speaker turns are normalized, and noise artifacts are removed.

The LLM processing stage forms the heart of the system. Carefully engineered prompts guide the LLM to generate three outputs: a concise meeting summary, a set of action items with assigned owners and deadlines, and a formal MoM (Laskar et al., 2023). Each of these outputs is generated in natural language but subsequently structured into a machine-readable JSON schema, enabling system interoperability.

The JSON outputs serve as inputs to external applications. Action items are automatically pushed to task trackers (e.g., Jira, Trello, Asana, or ClickUp) through their APIs (Asthana et al., 2025). Each task is tagged with metadata such as the meeting ID, deadline, and responsible person. Simultaneously, stakeholders receive personalized email notifications containing the meeting summary and their respective tasks. Email services like SendGrid or SMTP integration enable automated communication (Deng et al., 2023).

Completion tracking is realized in two modes. In the tracker-based mode, the system polls task trackers to update statuses in the database. In the email-based mode, stakeholder replies are parsed using the LLM to infer task completion (Sadia et al., 2025). For example, if a user responds, “I submitted the draft yesterday,” the model infers task completion and updates the database.

Finally, a centralized dashboard visualizes the current status of meetings, tasks, and completions. Filters allow managers to track performance by meeting ID, user, or deadline, while reports highlight overdue tasks and completion rates (Tan et al., 2023). The system emphasizes agentic behavior by not only recording outcomes but also proactively monitoring, reminding, and adapting to ensure accountability.

5. Benefits

The system provides several organizational benefits. First, it enhances time efficiency by automating labor-intensive processes of note-taking, summarization, and task allocation (Asthana et al., 2025).

Second, it ensures accuracy and consistency, producing standardized MoMs and reducing human error. Third, the agentic AI approach enables proactive reminders and accountability checks, thereby strengthening follow-through. Fourth, the system's integration with task trackers and dashboards fosters transparency and collaboration, enabling teams to stay aligned (Rennard et al., 2023). Finally, the design supports scalability, handling concurrent meetings across diverse teams (Li et al., 2019).

6. Challenges and Limitations

Despite its promise, the architecture faces several challenges. Speech recognition limitations include difficulties with unknown speakers, overlapping speech, and background noise (Schneider et al., 2025). Language barriers pose further issues, particularly in multilingual meetings where ASR systems may underperform. Task extraction challenges arise when tasks are ambiguously phrased or ownership is unclear (Golia & Kalita, 2023). LLMs also risk hallucination, where inaccurate or irrelevant content is generated (Tan et al., 2023). On a systems level, scalability may be constrained by computational demands, especially for real-time processing of multiple meetings. Ethical and security concerns include the need for role-based access, encryption, and compliance with organizational policies (Deng et al., 2023). Finally, over-reliance on automation could reduce human oversight, leading to blind acceptance of system outputs without critical review (Rennard et al., 2023).

7. Future Directions

Future research should address the limitations identified above. Enhancing speaker diarization and overlap detection can improve transcript accuracy (Schneider et al., 2025). Incorporating multimodal inputs, including video and contextual metadata, can enrich task extraction (Li et al., 2019). Expanding cross-lingual capabilities will support global teams (Sadia et al., 2025). Another direction lies in personalized AI agents capable of adapting to individual communication styles and task management preferences (Asthana et al., 2025). Integrating human-in-the-loop models will ensure that automation remains accountable and trustworthy (Rennard et al., 2023). Finally, establishing benchmarks for LLM accuracy in meeting contexts will enable standardized evaluation (Laskar et al., 2023).

8. Conclusion

This paper has presented a comprehensive LLM-centric architecture for an AI-powered meeting assistant, designed around the principles of agentic AI. By unifying speech recognition, LLM-driven summarization, structured JSON conversion, and workflow integration, the system automates the entire post-meeting lifecycle. The literature review highlighted advances in ASR, meeting summarization, task extraction, and workflow orchestration that underpin the design. The methodology demonstrates how speech inputs can be transformed into actionable outcomes seamlessly integrated with task trackers, notifications, and dashboards. While significant benefits are evident, challenges such as speech variability, ambiguous task phrasing, and ethical considerations remain. Addressing these will require ongoing research in agentic AI, multimodal systems, and human-AI collaboration. Nevertheless, the proposed system represents a major step toward transforming meetings from passive discussions into actionable, accountable workflows, unlocking new levels of organizational productivity.

Reference :

1. Schneider, F., Turchi, M., & Waibel, A. (2025). Policies and Evaluation for Online Meeting Summarization (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2502.03111

Vol-3 Issue-1 2025
Scientific Research Journal of Science, Engineering and Technology
ISSN: 2584-0584, Peer Reviewed Journal

2. Laskar, M. T. R., Fu, X.-Y., Chen, C., & TN, S. B. (2023). Building Real-World Meeting Summarization Systems using Large Language Models: A Practical Perspective (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2310.19233
3. Asthana, S., Hilleli, S., He, P., & Halfaker, A. (2025). Summaries, Highlights, and Action Items: Design, Implementation and Evaluation of an LLM-powered Meeting Recap System. Proceedings of the ACM on Human-Computer Interaction, 9(2), 1–29. https://doi.org/10.1145/3711074
4. Golia, L., & Kalita, J. (2023). Action-Item-Driven Summarization of Long Meeting Transcripts (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2312.17581
5. Jin, Y., Shi, Q., & Liu, Q. (2025). CFAS: consensus-focused abstractive meeting summarization through multi-party discourse modeling. Journal of King Saud University Computer and Information Sciences, 37(7). https://doi.org/10.1007/s44443-025-00210-3
6. Feng, X., Feng, X., Qin, B., & Geng, X. (2020). Dialogue Discourse-Aware Graph Model and Data Augmentation for Meeting Summarization (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2012.03502
7. Zhu, C., Xu, R., Zeng, M., & Huang, X. (2020). A Hierarchical Network for Abstractive Meeting Summarization with Cross-Domain Pretraining (Version 4). arXiv. https://doi.org/10.48550/ARXIV.2004.02016
8. Tan, H., Wu, H., Shao, W., Zhang, X., Zhan, M., Hou, Z., Liang, D., & Song, L. (2023). Reconstruct Before Summarize: An Efficient Two-Step Framework for Condensing and Summarizing Meeting Transcripts (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2305.07988
9. Deng, Z., Yoon, S., Bui, T., Dernoncourt, F., Tran, Q. H., Liu, S., Zhao, W., Zhang, T., Wang, Y., & Yu, P. S. (2023). Aspect-based Meeting Transcript Summarization: A Two-Stage Approach with Weak Supervision on Sentence Classification (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2311.04292
10. Rennard, V., Shang, G., Hunter, J., & Vazirgiannis, M. (2023). Abstractive Meeting Summarization: A Survey. Transactions of the Association for Computational Linguistics, 11, 861–884. https://doi.org/10.1162/tacl_a_00578
11. Li, M., Zhang, L., Ji, H., & Radke, R. J. (2019). Keep Meeting Summaries on Topic: Abstractive Multi-Modal Meeting Summarization. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics. https://doi.org/10.18653/v1/p19-1210
12. Sadia, B., Adeeba, F., Shams, S., & Hussain, S. (2025). Leveraging LLMs for action item identification in Urdu meetings: Dataset creation and comparative analysis. Information Processing & Management, 62(3), 104071. https://doi.org/10.1016/j.ipm.2025.104071