# AI-POWERED MEETING MINUTES GENERATOR USING BART & T5 (Hybrid Model)

**[1]G. Krishnaveni, [2]P. S. M. G. Sai Viswanath, [3]N. Sri Manohar Reddy, [4]S. Bhavani, [5]T. S. D. Satya Santhosh**

*[1]Assistant Professor, Department of Information Technology, Sir C R Reddy College of Engineering, Eluru*

*[2,3,4,5].Tech Students, Department of Information Technology, Sir C R Reddy College of Engineering, Eluru*

**ABSTRACT**

With an abundance of data around, the skill to effectively summarize important information from long conversations is more important than ever. This paper presents a novel AI-driven meeting minutes generator based on a hybrid model that benefits from the best of two recent transformer approaches, BART (Bidirectional and Auto-Regressive Transformer) and T5 (Text-to-Text Transfer Transformer). The system is optimized to automatically generate detailed meeting transcripts into short, context-sensitive summaries, thus decreasing the manual work of meeting documentation and improving productivity in general. Utilizing fine-tuned pre-trained models and state-of-the-art attention mechanisms, our system strives to capture critical discussion topics, decisions, and action items with high accuracy. Initial trials of comparable systems show encouraging performance in terms of workload saving and summary accuracy [1-2]. This study makes its own contribution to the increasing literature of automated meeting summarization through proposing and elaborating a hybrid model which should better the inadequacies of the current tools and models.

**Keywords:** Meeting minutes generation, Transformer Models, BART & T5, Attention Mechanism, ROUGE-L, CNN/DM, AMI Corpus, Automated Documentation, Keyword extraction.

## 1. INTRODUCTION

Meetings are the signature of the business life today, an indispensable tool of cooperation, decision-making, and information exchange. Yet, the usual practice of having one of the parties take meeting minutes manually is usually cumbersome, time-consuming, and error-prone, to the extent of possible miscommunication and lost efficiency. The sheer number of meetings multiplying in the busy business life today only makes the problems worse. The introduction of Artificial Intelligence (AI) and Natural Language Processing (NLP) technology promises unprecedented capacity to automate and improve the precision of meeting minuteing tasks.

Building upon the ground covered in prior work in automated text summarization [3], this document presents an AI-powered meeting minutes generator leveraging the synergy potential of BART and T5 fine-tuned transformer models [4]. These models have been seen to be incredibly effective in creating short and reasonable abstractions of humongous textual data [4]. Using these state-of-the-art models in a hybrid system, our system aims to automate the process of condensing lengthy meeting transcripts into short, contextually correct summaries, thereby minimizing human involvement and maximizing overall efficiency.

The proposed system is rendered scalable, secure, and user-friendly to accommodate diverse needs of organizations that are interested in efficient documentation systems. Preliminary analysis of AI-based approaches in this context has yielded notable reductions in admin burden as well as high summary accuracy [1-2]. This work tries to explore and further tune such AI-based approaches for accomplishing minute generation.

## 2. LITERATURE SURVEY

The area of automatic meeting minutes generation has received much prominence in recent years, and researchers have investigated numerous methods to tackle the challenges associated with summarizing verbal dialogue. In this paper, we give an overview of the related works that present the progress and various methods in this area:

Haz et al. (2024) examined open-source models for audio-to-text, summarization, and keyword extraction to design a modular structure to generate meeting minutes from captured presentations. Design centered on convenience of access as well as ease of modification for supporting future refinement. The system did well for audio-to-text transcription with minimal word-error rate, although there was potential to improve summarization and keyword extraction for precision as well as

relevance [5].
Matt Maxwell, AGB OnBoard (2024) introduced Automated Meeting Minutes feature that utilizes AI to make meeting minute creation easier. The feature records meetings, creates transcripts, and auto-generates minutes, making the process easier and lessening the workload for board professionals. This is one way that AI practicality and advantage come to the fore in meeting documentation [6].

Jiang (2023) focused on abstractive summarization of news articles with fine-tuned large language models, T5, PEGASUS, and BART. The work shares insights into their performance on news summarization tasks to compare their performance in creating high-quality summaries [7].

Yenigun (2025) provided a novice's manual for optimizing the T5 model for text summarization. The manual discusses data preparation, model training, and ROUGE metrics-based assessment, offering a step-by-step process for those seeking to improve text summarization methods [8].

Wijaya (2025) demonstrated how to summarize scientific papers using the BART model with the Hugging Face Transformers library. The paper explores text extraction, summarization, and hierarchical summarization techniques to improve the summarization of complex academic content [9].

## 3. REVIEW OF EXISTING MEETING SUMMARIZATION TOOLS

**Existing Tools and Their Limitations**
In the changing landscape of artificial intelligence (AI), several tools have been created to improve meeting productivity through automated transcription and summarization.
**Otter.ai:** Offers real-time transcription and collaborative note-taking, along with speaker labeling and keyword highlight. But does not handle nuances, context, and crosstalk well, often misclassifying sarcasm and requiring good-quality inputs in order to perform correctly. Also only supports a limited number of languages at the moment [10].
**MeetGeek:** Provides automated recording and transcription of meetings, with insights and summaries derived from AI. Although these capabilities are present, it might not be sophisticated enough to understand context completely, resulting in potential misinterpretation of context-rich discussions [11].
**IBM Watson:** Provides powerful speech-to-text functionality with variable language models that can be adapted. But in practice, usually needs to be heavily customized and requires technical experience, so not as easy for organizations without full-time technical support [12].
**Google Docs:** When integrated with AI add-ons, facilitates real-time collaboration and editing. Nevertheless, it lacks native advanced AI-driven transcription and summarization features, necessitating reliance on third-party integrations that may not seamlessly align with user needs [13].

## 4. PROPOSED SYSTEM

### 4.1 Overview of Our Models

In order to counter the shortcomings of present tools and models, we introduce a hybrid AI-based meeting minutes generator by fusing the best features of BART (Bidirectional and Auto-Regressive Transformers) and T5 (Text-to-Text Transfer Transformer) transformer models. This model would seek to utilize BART's ability in producing smooth and coherent abstractive summaries with T5's speed and performance in text-to-text applications, such as brief summarization. By integrating these models, we are hoping to get a system that generates accurate, coherent, and concise meeting minutes that can fit different meeting formats.

### 4.2 System Architecture

The proposed system architecture involves several key stages to process meeting data and generate the final minutes:

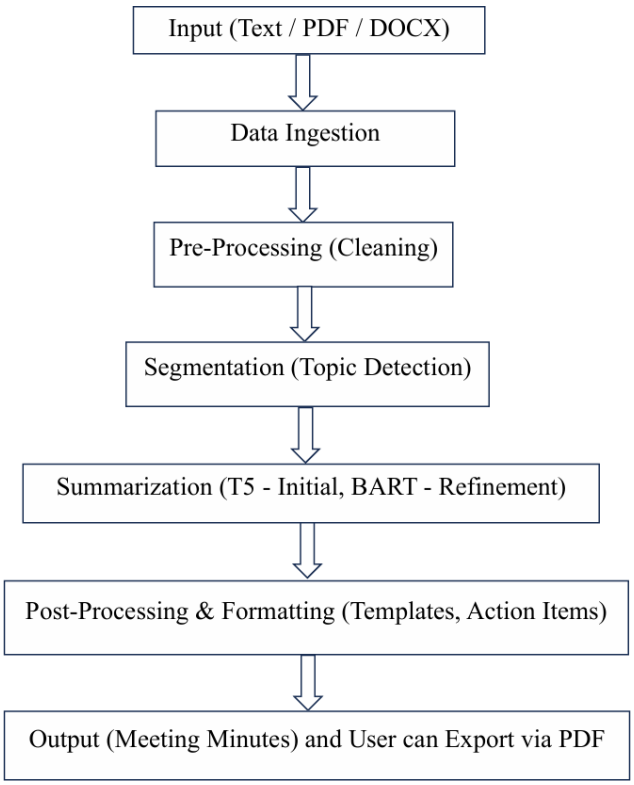


Fig 1: System Architecture

### 4.3 Software Design & Workflow

**The workflow of our AI-powered meeting minutes generator consists of the following steps:**

1) **Data Ingestion:** The system accepts meeting data in various forms like text transcripts, PDF documents, and DOCX documents [14].

2) **Preprocessing:** The input data is preprocessed to improve quality for use in summarization. It includes text cleaning by removing filler words and redundant information, speaker tagging to identify contributions differently, and organizing the content into a formal structure [15].

3) **Segmentation:** The pre-processed text is then divided into coherent, topic-specific segments to facilitate focused summarization [16]. Algorithms are employed to detect topic shifts and segment the text accordingly.

4) **Summarization:** The sentences are first summarized by the T5 model to produce structured summaries that capture the essence of the debates. The rough summaries are then fine-tuned by the BART model to enhance fluency, coherence, and clarity.

5) **Post-Processing and Formatting:** Post-processing summaries of all the sections are combined into one document and formatted according to pre-designed templates to get consistency and readability. Decisions, action items, and discussion points at this stage are also decided and formatted.

6) **Output (Meeting Minutes):** Users can export summarized output in PDF format

## 5. TRANSFORMER MODELS

### 5.1 About BART and T5

**1) BART (Bidirectional and Auto-Regressive Transformer):**

BART is a sequence-to-sequence model proposed by Facebook AI in 2019. It has the best of a bidirectional encoder (such as BERT) and an autoregressive decoder (such as GPT) [17]. BART proves to be excellent for natural language generation tasks like abstractive text summarization, machine translation, question answering, and text generation [18]. It can generate very brief summaries that extract the main essence of the source text. Its bidirectional encoder provides a comprehension of the complete context, while the autoregressive decoder makes it fluent and coherent.

**Working Mechanism**

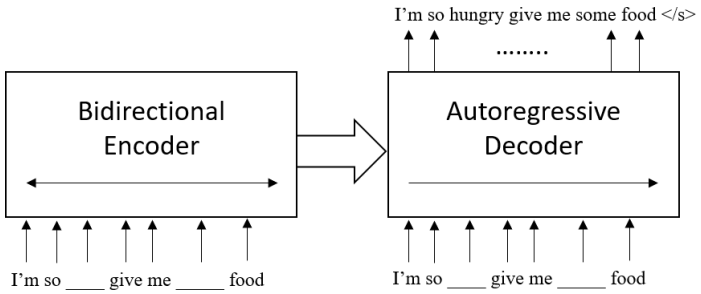


Fig 2: Encoder & Decoder of BART

**Encoder:** The encoder scans the entire input text in both directions such that it has access to the context of each token in terms of the overall input. This helps in extracting subtle relationships among words [17].

**Decoder:** The decoder generates the output (e.g., summary) autoregressively, token by token, using previously generated tokens as context [17].

**Pre-training:** BART is pre-trained on a denoising autoencoder. It discovers how to transform noisy versions back into the initial text using diverse noising mechanisms like token masking, deletion, text infilling, sentence rotation, and document rotation [20].

**Fine-tuning:** Then, BART can be fine-tuned from specific downstream tasks like text summarization using task-specific datasets [19,20].

**2) T5 (Text-to-Text Transfer Transformer):**

T5, or Text-to-Text Transfer Transformer, is another popular transformer-based language model introduced by Google Research in 2019. It follows an encoder-decoder structure inspired by the base Transformer model. It takes a distinct "text-to-text" framework, viewing all NLP tasks like summarization, translation, question answering, and classification as text input to text output conversion. This single framework enables T5 to process various tasks with the same architecture, loss function, and hyperparameters. [21,23]

T5 is trained beforehand on the Colossal Clean Crawled Corpus (C4), a large collection of web-crawled text and code, T5 acquires general language understanding and generation skills, which are fine-tuned to address specific applications [21]. The framework also renders T5 ready for structured text generation workloads, e.g., meeting minutes summarization, where it can produce particular subsections such as action items and decisions.

**Working Mechanism**

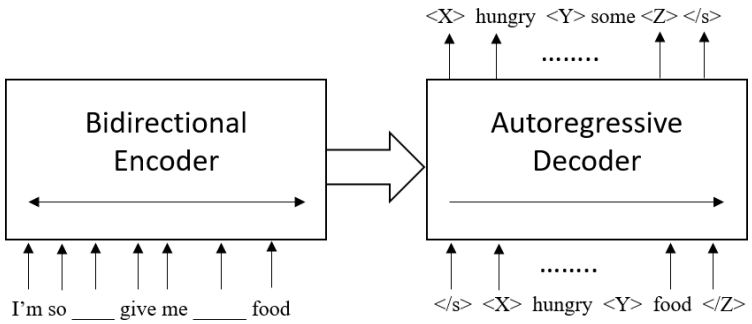

Fig 3: T5 only outputs the masked parts that need to be predicted

**Encoder-Decoder Architecture:** T5 uses a simple Transformer encoder-decoder architecture [21].
**Pre-training:** T5 is pre-trained on the vast C4 (Colossal Clean Crawled Corpus) corpus. The pre-training task involves masked language modeling, where the model learns to predict removed words in a text [24,25]
**Task Prefixing:** For several tasks, the input text is typically prefixed with an instruction. For summarization, the input is typically prefixed with "summarize:".
**Fine-tuning:** T5 can be fine-tuned effectively on specific datasets for several NLP tasks, including summarization [24].

### 5.2 How Our Proposed Hybrid Model (T5 -> BART) Overcomes Challenges:

- **Phase 1 (T5):** We use T5's text-to-text speed and agility to generate an initial draft of meeting minutes. By using clear instructions like "summarize: include action items and deadlines," we can ask T5 to pull out useful information and even attempt to find action items and deadlines. T5's speed allows for rapid generation of this initial draft [26].
- **Phase 2 (BART):** T5's first draft is then handed over to BART for polishing. BART's context retention and coherence strength are utilized here. BART can enhance the logical sequence of the summary, correct any probable hallucinations or errors, and refine the natural language fluency of the generated minutes. BART's capacity to cope with noisy inputs (such as possible disfluencies in T5's first draft) also proves useful [27,28].

This sequential flow leverages T5's speed and conciseness for information extraction and BART's strength in natural language generation for producing high-quality minutes.

### 5.3 Strengths

- **BART Strengths:** Pros at creating fluent and coherent abstractive summaries, robust to noise, is able to generate grammatically correct text and has a satisfactory quality-efficiency trade-off [27,28]
- **T5 Strengths:** Robust text-to-text model, produces good short abstractions, does well on a wide range of NLP tasks, and can be finely tuned well [22,26].
- **Hybrid Model Strengths:** It integrates the fluency and coherence of BART with the efficiency and brevity of T5, potentially resulting in more detailed and user-friendly meeting minutes.

## 6. RESULTS

### 1) User Interface (React UI)

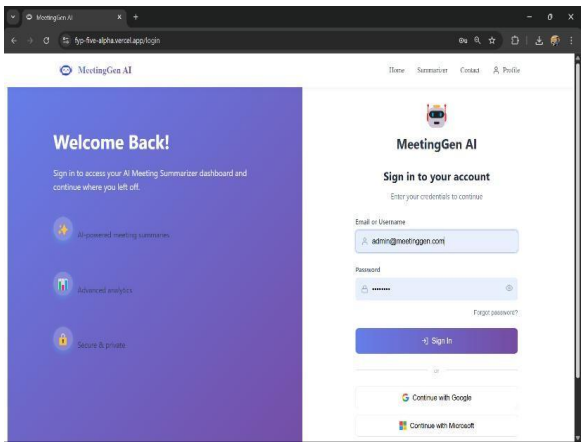


Fig 4: User Authentication via Login Page

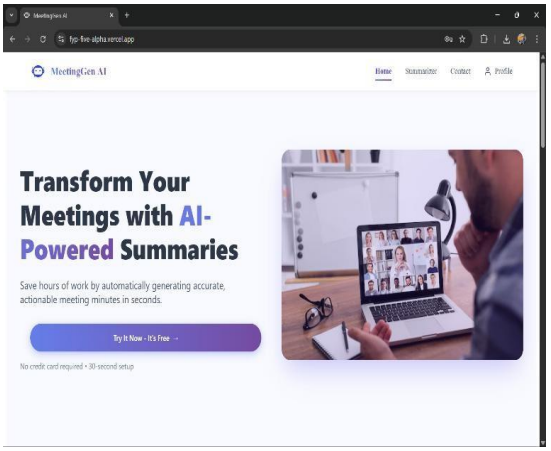


Fig 5: Home Page

MeetingGen AI platform provides users with a convenient and secure login process, wherein they can use valid credentials or Google and Microsoft accounts to log in. The clean navigation and proper call-to-action assist the users in browsing features like summarization, contact, and profile management.

### 2) User Request Handling: Text and File Submission

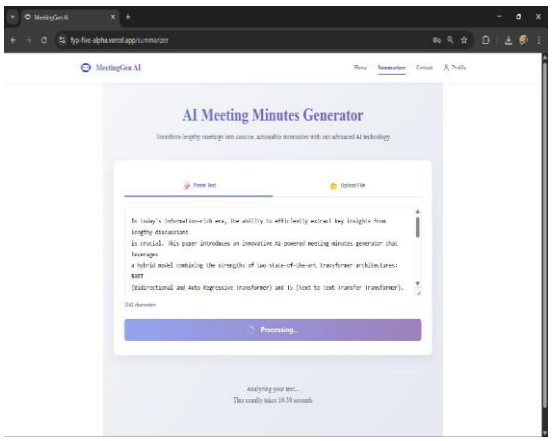


Fig-6: User Request in Text Format

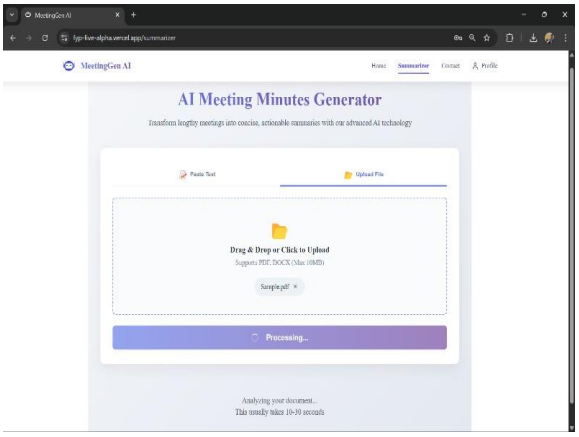


Fig-7: User Request in DOCX or PDF Format

Summarizer Page of MeetingGen AI has two modes of input for processing transcripts of a meeting: text mode and file upload. In text mode, users copy and paste their transcript and click on "Generate Summary," which initiates a loading state with visible progress feedback and live character count. In file upload mode, users upload files (PDF/DOCX, maximum 10MB), where the uploaded file is shown and same loading animation is displayed. Both modes display a message indicating the AI is processing the input, which allows for a good and interactive user experience while processing. Request Received by the Server – Status 200 OK

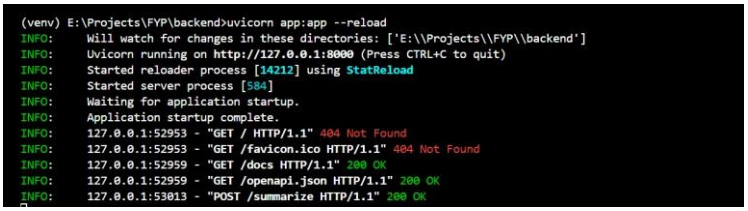


Fig-8: Request Received by the Server – Status 200 OK

This picture shows the server's reply that the request has been accepted and processed successfully. The HTTP status code of 200 OK means that the system is running properly and in a position to produce the summary.

### 3) System Processing Workflow: Text and File Handling

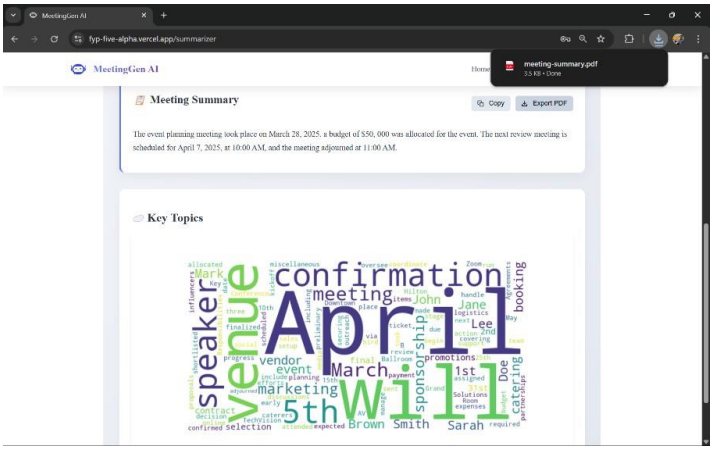


Fig-9: Summarized Output with Word Cloud

On processing, MeetingGen AI displays two sections: a formal Meeting Summary and Key Topics word cloud. The summary provides vital information such as date, significant decisions made, future meeting timing, and adjournment time with copy or export as PDF options. The lower word cloud reads most frequent terms appearing, providing a visual snapshot of discussion main points at a glance. This output screen finishes summarization in neat, actionable information.

## 7. CONCLUSION

This paper suggests an AI-aided meeting minutes generator using a hybrid BART-T5 model to facilitate automation of converting meeting transcripts into coherent and useful meeting minutes with minimal human effort and maximum productivity gain. The new architecture is that of a systematic workflow of data ingestion, preprocessing, segmentation, summarization with the hybrid BART-T5 model, and post-processing formatting with allowances for preserving context, speaker-aware summarization, and structured output generation. The advantages of our system over existing tools, along with its possibilities for use by various groups of users, indicate its potential for change in this study. Additional improvement will also include real-time summarization to create automatically meeting minutes while meetings are on-going, multi-language support for transcription and summarization across different languages, and summary template customizability so that users can generate meeting minutes based on specific requirements. With such developments, the AI-Powered Meeting Minutes Generator with BART and T5 (Hybrid Model) will be an even more adaptive, scalable, and intelligent solution for organizations and experts who require automated, high-quality meeting minutes.

## 8. REFERENCES

[1]. https://arxiv.org/html/2312.17581v2
[2]. https://www.ibm.com/think/insights/generative-ai-benefits
[3]. https://github.com/carlosrod723/BART-Abstract-Text-Summarization
[4]. https://stackoverflow.com/questions/77732511/what-are-differences-between-t5-and-bart
[5]. Haz, Amma & Panduman, Yohanes & Funabiki, Nobuo & Fajrianti, Evianita. (2024). Fully Open-Source Meeting Minutes Generation Tool. Future Internet. 16. 429. 10.3390/fi16110429.
[6]. https://agb.org/blog-post/the-ai-powered-meeting-minutes-tool-is-here
[7]. https://www.kaggle.com/code/zifengjiang/news-summarization-t5-pegasus-and-bart
[8]. https://pub.aimind.so/fine-tuning-t5-for-summarization-a-beginners-guide-1d0fce60f680
[9]. https://www.kdnuggets.com/how-to-summarize-scientific-papers-bart-hugging-face-transformers
[10]. https://www.bitdegree.org/ai/otter-ai-review
[11]. https://meetgeek.ai/ai-meeting-minutes
[12]. https://www.ibm.com/architectures/hybrid/genai-document-summarization
[13]. https://fellow.app/blog/productivity/google-meet-gemini-ai-note-taker-review/
[14]. https://medium.com/@rehabreda/build-an-ai-meeting-minutes-generator-using-whisper-and-deepseek-llm-for- automated-meeting-c5345d392e75
[15]. https://arxiv.org/html/2407.11919v1
[16]. https://arxiv.org/html/2312.17581v2
[17]. https://medium.com/@risharane6/bart-transformer-for-text-summarization-by-risha-rane-739ffbf5af30
[18]. https://www.projectpro.io/article/transformers-bart-model-explained/553
[19]. https://www.analyticsvidhya.com/blog/2024/11/bart-model/
[20]. https://www.digitalocean.com/community/tutorials/bart-model-for-text-summarization-part1
[21]. https://medium.com/@gagangupta_82781/understanding-the-t5-model-a-comprehensive-guide-b4d5c02c234b
[22]. https://www.analyticsvidhya.com/blog/2024/05/text-summarization-using-googles-t5-base/
[23]. https://medium.com/artificialis/t5-for-text-summarization-in-7-lines-of-code-b665c9e40771
[24]. https://medium.com/@lmpo/encoder-decoder-transformer-models-a-comprehensive-study-of-bart-and-t5- 132b3f9836ed
[25]. GUNDA, P., & Thirupathi Rao KOMATI. (2024). Integrating Self-Attention Mechanisms For Contextually Relevant Information In Product Management. International Journal of Computational and Experimental Science and Engineering, 10(4). https://doi.org/10.22399/ijcesen.651
[26]. https://www.restack.io/p/ai-summarization-answer-t5-model-cat-ai
[27]. https://medium.com/@gabya06/t5-vs-bart-the-battle-of-the-summarization-models-c1e6d37e56ca
[28]. https://www.width.ai/post/bart-text-summarization