

MEDSPEAK: A KNOWLEDGE GRAPH-AIDED ASR ERROR CORRECTION FRAMEWORK FOR SPOKEN MEDICAL QA

Yutong Song^{1*}, Shiva Shrestha², Chenhan Lyu¹, Elahe Khatibi¹,
Pengfei Zhang¹, Honghui Xu², Nikil Dutt¹, Amir Ramani¹

¹University of California, Irvine, Computer Science, Irvine, CA, USA

²Kennesaw State University, Computer Science, Kennesaw, GA, USA

ABSTRACT

Spoken question-answering (SQA) systems relying on automatic speech recognition (ASR) often struggle with accurately recognizing medical terminology. To this end, we propose MedSpeak, a novel knowledge graph-aided ASR error correction framework that refines noisy transcripts and improves downstream answer prediction by leveraging both semantic relationships and phonetic information encoded in a medical knowledge graph, together with the reasoning power of LLMs. Comprehensive experimental results on benchmarks demonstrate that MedSpeak significantly improves the accuracy of medical term recognition and overall medical SQA performance, establishing MedSpeak as a state-of-the-art solution for medical SQA. Our codes are available in <https://github.com/RainieLLM/MedSpeak>.

Index Terms— Knowledge Graph, Automatic Speech Recognition, Spoken Question-Answering, LLM

1. INTRODUCTION

The increasing adoption of speech-based AI systems in medical applications has led to the growing need for high-accuracy spoken medical question answering (SQA) systems. These systems rely on Automatic Speech Recognition (ASR) to transcribe spoken medical queries before feeding them into retrieval-augmented generation (RAG) models or large language models (LLMs) for answer generation. However, despite advancements in ASR models, such as Whisper and Wav2Vec2, their performance remains suboptimal in medical settings due to several domain-specific challenges [1, 2]. These challenges include complex medical terminology, phonetic ambiguities, and contextual dependencies, which degrade ASR transcription accuracy and impact downstream medical reasoning [3, 4]. In other words, traditional ASR models, trained on general-domain corpora, often struggle to accurately recognize specialized medical terminology, leading to frequent misrecognitions of critical medical entities [5, 6]. Although some domain-specific finetuning approaches have been attempted [2, 7, 8], they are restricted by

data insufficiency and lack of generalization. Additionally, these domain-specific models still face challenges with phonetic confusability between terms like "hypertension" and "hypotension," further exacerbating ASR errors [9]. While distinct in their medical implications, those words can be confusingly similar to ASR systems, leading to errors in transcription. These deficiencies have negative impacts in clinical applications where transcription errors cascade to incorrect reasoning with resulting threats to patient safety. This underscores the need for a robust ASR with a reliable QA system to support trustworthy healthcare decisions.

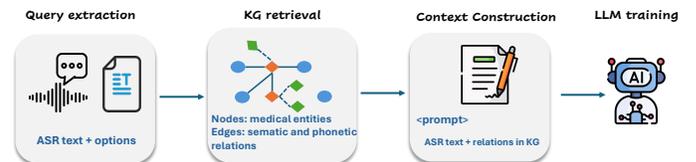


Fig. 1. Overview of MedSpeak.

To address these challenges, in this paper, we propose MedSpeak, a knowledge graph-aided ASR error correction framework designed specifically for spoken medical QA tasks. By harnessing the power of knowledge graphs and the reasoning capabilities of LLMs, MedSpeak integrates structured domain knowledge, phonetic and semantic similarity constraints, and context-aware reasoning to significantly improve medical ASR correction and QA ability, as shown in Fig. 1. Our contributions include: (1) An automated medical knowledge graph (KG) construction method encodes semantic and phonetic relationships between terms, enabling robust error correction and reasoning. (2) A multi-constraint retrieval mechanism leverages phonetic and semantic features to generate accurate ASR hypotheses for medical terminology. (3) A fine-tuned LLM integrates retrieved knowledge with answer constraints, producing precise transcriptions and reliable answers. (4) Extensive experiments on spoken medical QA datasets show that MedSpeak achieves state-of-the-art ASR error correction in the medical domain.

*Corresponding author: yutons12@uci.edu

2. RELATED WORKS

Traditional ASR models, trained on general-domain corpora, often struggle to accurately recognize specialized medical terminology, leading to frequent misrecognitions of critical medical entities [1, 5, 6, 7]. While fine-tuning ASR models on domain-specific datasets has been explored, these efforts are still limited by data scarcity and poor generalization [2, 7, 8]. Additionally, these models face challenges with phonetic confusability between terms like "hypertension" and "hypotension," further exacerbating ASR errors [3, 4, 9].

To address these challenges, RAG-based methods aim to improve ASR accuracy and named entity recognition [10] by incorporating external domain specific knowledge retrieval [11, 12, 13]. Techniques such as GEC-RAG [2] and IRAG [14] utilize contextually retrieved text snippets to refine ASR hypotheses. However, these approaches relies on text-centric evidences and fail to leverage phonetic knowledge, limiting their ability to resolve confusion between phonetically similar medical entities. Structured knowledge graphs have been explored for domain-specific ASR correction, particularly in medical QA systems. MedRAG [15] incorporates a medical KG into retrieval-based ASR correction but does not explicitly address phonetic similarity for disambiguation. Similarly, RANEC [3] leverages named entity correction for ASR refinement but overlooks the integration of multiple-choice context for better accuracy.

3. MEDSPEAK

Our MedSpeak framework combines static KG context injection with a fine-tuned LLM, explicitly integrating both semantic relationship and phonetic similarity into LLM fine-tuning. This combined representation allows for robust ASR error correction and QA reasoning within a unified two-line supervision format, achieving great performance across multiple medical QA benchmarks. Fig. 2 presents an example

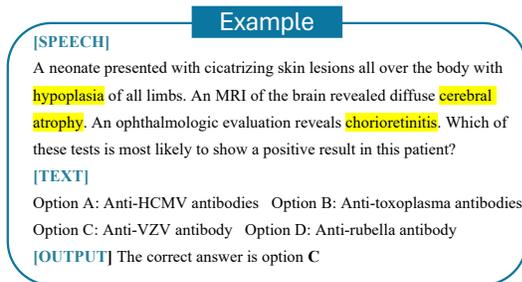


Fig. 2. An Example of SQA.

of our spoken question-answering (SQA) framework. The system receives two distinct inputs: (1) a clinical case description in speech form (containing key diagnostic terms highlighted in yellow such as "hypoplasia," "cerebral atrophy," and "chorioretinitis"), and (2) four possible diagnostic test options provided as text (Options A-D). The system then produces two outputs: the accurately transcribed speech and

the correct option (Option C). This example demonstrates the critical challenge of accurately recognizing specialized medical terminology in speech, as mis-recognition of terms like "chorioretinitis" as "chorioamnionitis" would lead to incorrect evidence retrieval and subsequent diagnostic errors. Our proposed MedSpeak leverages knowledge graph-aided error correction to improve recognition of these critical diagnostic terms, ensuring that the speech-to-text conversion preserves the medical meaning necessary for selecting the appropriate diagnostic test. In the following, we detail two main components in MedSpeak.

3.1. Knowledge Graph Building

We construct our knowledge graph using the Unified Medical Language System (UMLS) dataset from the National Institutes of Health (NIH) [16], establishing relationships between medical terms through the MRREL table. During construction, we iteratively select medical keywords from the table while filtering out generic relationships, duplicate entries, and relationship pairs that lack specific semantic information. After each keyword is selected, we use The CMU Pronouncing Dictionary to generate phonetically similar words and append them to our knowledge graph. By removing these uninformative relationships and adding pronunciation information, we streamline both the knowledge graph construction process and improve retrieval efficiency. As outlined in Fig. 1, in the "Relations in the Knowledge

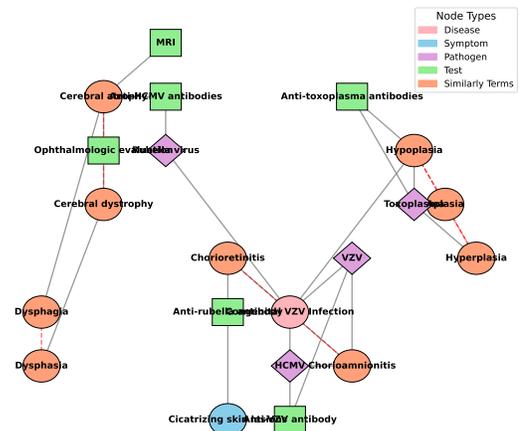


Fig. 3. Medical KG Constructed from SQA.

Graph" step, we deploy the Double Metaphone algorithm and Levenshtein Distance metric to identify English words with similar pronunciation to medical entities. Since Double Metaphone alone produces excessive false positives, combining both metrics enhances phonetic matching accuracy. We then integrate these entities into a unified knowledge graph for retrieval. Fig.3 depicts a semantic knowledge graph for the QA example in Fig.2. Each node represents differential diagnostic considerations (HCMV, Toxoplasma, VZV, and Rubella virus infections) with corresponding antibody

tests. The graph incorporates phonetically similar terms often misrecognized in speech systems, such as "Hypoplasia"/"Hyperplasia" and "atrophy"/"hypertrophy." In the final KG, nodes contain unique identifiers and term strings, while edges have relationship tags including "classifies," "constitutes," "due_to," and "plays_role" to help the LLM understand relationships between question terms. A specific "phonetic" edge tag aids the LLM in identifying potential speech transcription errors.

3.2. KG-Aided ASR Error Correction Framework

In our MedSpeak, the workflow begins with an existing speech recognition model to generate transcripts, followed by identification of medical terminology for knowledge graph retrieval. The retrieved medical information is then embedded alongside contextual prompts to enhance LLM reasoning and correction capabilities.

3.2.1. Training Process

Our training methodology primarily fine-tunes the parameters of a base LLM model under the full model supervision. Each training sample is formatted as a dialogue with following three parts: **i) System Messages** consists of a fixed rules only instruction that enforced a strict two-line output from the LLM. **ii) User Input** included three items:(a) the noisy ASR transcript, (b) the multiple-choice answer options, and (c) the semantic and phonetic Knowledge-Graph (KG). The KG section is truncated to a fixed budget(900 tokens) in order to prevent gold labels from being dropped. Finally, **iii) Assistant Target** is defined as a structured two-line completion format:

```
Corrected Text: [Original GT Text]
Correct Option: [Option Letter<A|B|C|D>]
```

Training Objective: Given an input-output pair (x, y) , the model is optimized with a single causal language modeling objective. Formally, the input sequence is defined as $x = [s; u(\text{ASR}, \text{Options}, \text{KG})]$, where s is the fixed *system instruction*, and $u(\cdot)$ represents the *user input* consisting of the ASR transcript, multiple-choice options, and truncated KG context. The target sequence is expressed as $y = \text{Corrected Text: } t^* \oplus \text{Correct Option: } o^*$, where t^* is the *gold-standard transcript (ground truth)* and $o^* \in \{A, B, C, D\}$ is the *correct answer option*. The model parameters θ are trained to minimize

$$\mathcal{L}(\theta) = - \sum_{i=1}^{|y|} \log P_{\theta}(y_i | x, y_{<i}), \quad (1)$$

where y_i is the i -th token in the target sequence, $y_{<i}$ are all previous tokens, and $P_{\theta}(y_i | x, y_{<i})$ is the probability of predicting the correct token given the input x and prior context. This formulation unifies *ASR error correction* (via t^*)

and *multiple-choice reasoning* (via o^*) into a single generation task. The KG ensures domain knowledge remains available, enabling the model to produce clinically accurate and context-aware outputs.

3.2.2. Inference Process

The acoustic input a is transcribed into a noisy text by the pre-trained ASR model which is given as $\hat{t} = \text{ASR}(a)$. The user input is then constructed in the same format as during training: $u = u(\hat{t}, \mathcal{O}, \mathcal{K})$, where \hat{t} is the ASR transcript, $\mathcal{O} = \{A, B, C, D\}$ denotes the multiple-choice options, and \mathcal{K} is the KG context $(\mathcal{K}_{\text{sem}}, \mathcal{K}_{\text{phon}})$. The overall model input is given by $x = [s; u]$, where s is the fixed system instruction. The fine-tuned LLM generates the structured two-line output:

$$y = \text{LLM}_{\theta}(x) \Rightarrow \begin{cases} \text{Corrected Text: } \tilde{t} \\ \text{Correct Option: } o^* \end{cases} \quad (2)$$

with \tilde{t} denoting the corrected transcript and $o^* \in \{A, B, C, D\}$ the predicted option.

4. EXPERIMENT AND RESULTS

In this section, we first describe the data preparation and experimental settings. We then evaluate our MedSpeak framework on spoken medical QA against various baselines.

4.1. Data Preparation

To systematically evaluate MedSpeak, we use a diverse medical SQA benchmark by synthesizing spoken data from three well-established multiple-choice QA datasets [1, 6, 17]. The MMLU Medical Tasks dataset includes 1,089 questions covering six medical subjects, including biology, anatomy, and clinical medicine. The MedQA (USMLE-Based) dataset consists of 1,273 questions derived from the United States Medical Licensing Examination (USMLE), incorporating real-world diagnostic reasoning tasks. The MedMCQA dataset comprises 4,183 questions spanning multi-specialty medical assessments used in Indian medical entrance exams. Each dataset is synthesized at 16,000 Hz, mono WAV format using local pyttts3 text-to-speech(TTS) engine for speech generation. This benchmark provides over 47 hours of spoken medical QA data, ensuring rigorous evaluation of ASR correction, retrieval-driven reasoning, and knowledge integration.

4.2. Experimental Settings

Configurations: Our MedSpeak framework is implemented by fine-tuning all parameters of Llama-3.1-8B-Instruct. We conduct training across approximately $\sim 10\text{k}$ of vocalized medical QA pairs drawn from MedMCQA, MedQA, and MMLU-Medical, totaling about 6.79M tokens. The model is optimized for 10 epochs with a batch size of 8 and gradient the building up of 16, using a learning rate multiplier of 2 relative to the cosine schedule baseline. The length of sequence is constrained to at most 2048 tokens, To prevent truncation of gold labels, we budget semantic and phonetic knowledge-graph partitions of 600 and 300 tokens respectively. The

objective of training is the standard causal language modeling loss across the whole rendered dialogue, so that the model jointly learns ASR correction and answer prediction. This experiment is conducted on a cluster of 8x NVIDIA A100 GPUs with 80 GB memory each, paired with AMD EPYC 7742 64-Core Processor. We enable `bfloat16` precision, gradient checkpointing, and TF32 matrix multiplication for efficiency enhancement. The training runs are distributed across multiple GPUs with the help of the HuggingFace Trainer framework. The Whisper Small Model (244M parameters) is used as the front-end ASR system to generate the noisy transcript for all our experiments.

Baselines: We compare our MedSpeak framework against a set of baseline systems to evaluate the contribution of fine-tuning, ASR correction and KG integration: (1) Zero-shot ASR (**ZS-ASR**) is the model using the raw whisper generated transcript without fine-tuning or KG-context. (2) Zero-shot GT (**Zero-Shot**) is the model using the pure base LLM when provided with the ground truth transcripts, and without fine-tuning or KG-context. (3) Fine-Tuned LLM + Whisper (**FT+Whisp**) is the model where the fine-tuned LLM is given the whisper generated transcripts without KG-context. (4) Fine-Tuned LLM + GT (**FT-LLM**) is the fully fine-tuned LLM evaluated with the ground truth texts. (5) **MedSpeak**: is our proposed model that incorporates noisy transcripts generated by whisper with the budgeted KG context and the fine-tuned LLM.

Evaluation Metrics: We evaluate MedSpeak performance on two metrics: (1) QA Accuracy (%): the fraction of correctly predicted multiple choice answers, and (2) Word Error Rate (WER) (%): Punctuation-insensitive Levenshtein distance between reference and corrected transcripts, which quantifies the ASR correction quality, calculated as: $WER = \frac{S+D+I}{N} \times 100\%$, where S are substitutions, D deletions, I insertions, and N is the total number of words in the reference transcript.

4.3. Results and Analysis

We compare the QA accuracy and WER of our proposed framework MedSpeak against different baselines for different datasets. Accuracy results are listed in Table 1. Indicatively, the best performance is recorded by MedSpeak overall on the benchmarks, recording an overall accuracy of 93.4%, very close to the performance of the FT-LLM baseline paired with the ground-truth texts (92.5%) and considerably higher than zero-shot setups in the 50-56% range. The biggest boosts are recorded in Virology and in MedQA, where MedSpeak reduces the gap by over four percentage points compared to the FT-LLM baseline, demonstrating our knowledge graph context integration to be highly beneficial in coping with challenging terminology and noisy ASR transcriptions. The exception is in the case of MedMCQA where the LLM fine-tuned on ground-truth transcriptions performs by a very slim margin (92.5% vs. 91.5%). Even then though, MedSpeak outperforms baselines in all other tasks. Moreover, WER

Table 1. QA Accuracy of MedSpeak and Baselines.

Task	ZS-ASR	Zero-Shot	FT+Whisp	FT-LLM	MedSpeak
MMLU					
Clinical	62.2	66.3	85.6	94.3	95.4
Anatomy	57.0	64.4	85.2	94.1	93.3
College	62.1	67.5	84.2	92.7	95.6
Virology	47.6	48.8	85.5	91.0	95.8
Prof. Med.	74.6	76.1	88.6	94.9	97.8
MedQA	54.9	58.9	86.6	91.8	97.5
MedMCQA	45.5	52.8	82.3	92.5	91.5
Avg.	50.2	56.3	83.7	92.5	93.4

results are presented in Table 2. As WER is defined only over ASR-processed inputs, ground-truth baselines are omitted; if used as input, their error rate would necessarily be zero. MedSpeak shows outright benefits. The overall WER decreases to 29.9%, versus 77.2% on ZS-ASR generated transcripts with non finetuned LLM and 35.8% on FT+Whisp. Gains are particularly significant in Anatomy and Virology, where MedSpeak decreases error by 5–7 percentage points over the next-best baseline. These decreases are significant to how well combining semantic and phonetic knowledge graph information with fine-tuning helps the model to discriminate phonetically similar but semantically distinct medical entities and to correct domain-specific transcription mistakes with greater reliability. These results in tandem confirm that MedSpeak is equally very accurate on spoken medical QA and improves transcription quality.

Table 2. WER of MedSpeak and Baselines.

Task	ZS-ASR	FT+Whisp	MedSpeak
MMLU			
Clinical	76.7	27.3	22.1
Anatomy	87.0	25.8	19.8
College	67.5	21.0	15.7
Virology	78.9	24.6	17.5
Prof. Med.	61.9	39.3	37.8
MedQA	63.9	44.9	43.5
MedMCQA	81.4	35.4	28.1
Avg.	77.2	35.8	29.9

5. CONCLUSION

In this paper, we proposed MedSpeak, a KG-aided ASR error correction framework designed for spoken medical question answering. The framework integrates semantic and phonetic context from a medical knowledge graph with the reasoning ability of LLMs, enabling robust correction of transcription errors and reliable clinical reasoning. A multi-constraint retrieval process was developed to align phonetic similarity and semantic relations with ASR hypotheses, while a fine-tuned LLM unifies error correction and answer generation within a single objective. Through extensive evaluation, MedSpeak demonstrated consistent improvements in both transcription accuracy and downstream QA reliability, offering a scalable and effective solution for enhancing ASR systems in high-stakes medical applications.

6. REFERENCES

- [1] Yanis Labrak, Adel Moumen, Richard Dufour, and Mickael Rouvier, “Zero-shot end-to-end spoken question answering in medical domain,” *arXiv preprint*, 2025.
- [2] Amin Robotian, Mohammad Hajipour, Mohammad Reza Peyghan, Fatemeh Rajabi, Sajjad Amini, Shahrokh Ghaemmaghami, and Iman Gholampour, “Gec-rag: Improving generative error correction via retrieval-augmented generation for automatic speech recognition systems,” *arXiv preprint*, 2025.
- [3] Ernest Pusateri, Anmol Walia, Anirudh Kashi, Bortik Bandyopadhyay, Nadia Hyder, Sayantan Mahinder, Raviteja Anantha, Daben Liu, and Sashank Gondala, “Retrieval-augmented correction of named entity speech recognition errors,” *IEEE*, 2025.
- [4] Mingqiu Wang, Izhak Shafran, Hagen Soltau, Wei Han, Yuan Cao, Dian Yu, and Laurent El Shafey, “Retrieval-augmented end-to-end spoken dialog models,” *ICASSP*, 2025.
- [5] Do June Min, Karel Mundnich, Andy Lapastora, Erfan Soltanmohammadi, Srikanth Ronanki, and Kyu Han, “Speech retrieval-augmented generation without automatic speech recognition,” *ICASSP*, 2025.
- [6] Xuejiao Zhao, Siyan Liu, Su-Yin Yang, and Chunyan Miao, “Medrag: Enhancing retrieval-augmented generation with knowledge graph-elicited reasoning for healthcare copilot,” *arXiv preprint*, 2025.
- [7] Hao Yang, Min Zhang, Minghan Wang, and Jiabin Guo, “Rasu: Retrieval-augmented speech understanding through generative modeling,” *IEEE*, 2025.
- [8] Chunyu Sun, Bingyu Liu, Zhichao Cui, Anbin Qi, Tianhao Zhang, Dinghao Zhou, and Lewei Lu, “Seal: Speech embedding alignment learning for speech large language model with retrieval-augmented generation,” *arXiv preprint*, 2025.
- [9] Yi Yuan, Haohe Liu, Xubo Liu, Qiushi Huang, Mark D. Plumbley, and Wenwu Wang, “Retrieval-augmented text-to-audio generation,” *ICASSP*, 2025.
- [10] Zhihong Lei, Xingyu Na, Mingbin Xu, Ernest Pusateri, Christophe Van Gysel, Yuanyuan Zhang, Shiyi Han, and Zhen Huang, “Contextualization of asr with llm using phonetic retrieval-based augmentation,” *arXiv preprint arXiv:2409.15353*, 2024.
- [11] Hao Yang, Min Zhang, Minghan Wang, and Jiabin Guo, “Rasu: Retrieval augmented speech understanding through generative modeling,” in *Interspeech*, 2024, vol. 2024, pp. 3510–3514.
- [12] Jinlong Xue, Yayue Deng, Yingming Gao, and Ya Li, “Retrieval augmented generation in prompt-based text-to-speech synthesis with context-aware contrastive language-audio pretraining,” *arXiv preprint arXiv:2406.03714*, 2024.
- [13] Cihan Xiao, Zejiang Hou, Daniel Garcia-Romero, and Kyu J Han, “Contextual asr with retrieval augmented large language model,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [14] Hao Yang, Min Zhang, and Daimeng Wei, “Irag: Iterative retrieval-augmented generation for spoken language understanding,” *IEEE*, 2025.
- [15] Xuejiao Zhao, Siyan Liu, Su-Yin Yang, and Chunyan Miao, “Medrag: Enhancing retrieval-augmented generation with knowledge graph-elicited reasoning for healthcare copilot,” *arXiv preprint*, 2025.
- [16] National Library of Medicine, “Unified medical language system (umls),” National Institutes of Health, 2024.
- [17] Rui Liu, Zhenqi Jia, Feilong Bao, and Haizhou Li, “Retrieval-augmented dialogue knowledge aggregation for expressive conversational speech synthesis,” *Information Fusion*, 2025.