Speech-to-Text and Audio Cleansing with AI: Improving Transcription Accuracy

Dipti Suryavanshi[1] Mr. Shripad Bhide[2]
[1]MCA Student [2]Assistant Professor
[1,2]Master Of Computer Applications
[1,2]PES Modern College of Engineering, Pune, India

Abstract — Speech-to-text technology has become a cornerstone in many areas such as virtual assistants, automatic transcription, accessibility tools, and real-time communication solutions. Our growing reliance on voice interactions has sparked significant advancements in AI-powered speech recognition systems. However, achieving high transcription accuracy presents a real challenge, especially with external factors like background noise, varying accents, crowded environments, and subpar audio quality. The rise of deep learning and AI-enhanced audio cleaning techniques—like noise reduction, echo cancellation, and speaker diarization—has notably improved transcription accuracy. Even so, we're still not quite reaching that human-like transcription quality, particularly in noisy settings or with low-resource languages. This paper takes a deep dive into AI-driven speech-to-text (STT) systems, exploring their evolution and impact on transcription accuracy. We compare various AI-driven STT models, such as OpenAI Whisper, Mozilla DeepSpeech, IBM Watson Speech-to-Text, and Google Speech-to-Text API, analyzing how they perform in different environments. By testing these models under various acoustic conditions, we assess their capabilities against the real-world problems they might face. We evaluate transcription quality using key metrics like Word Error Rate (WER), Signal-to-Noise Ratio (SNR), and latency issues. Additionally, we look into AI-based audio preprocessing methods to see how they help reduce errors and boost transcription quality. This includes examining deep learning noise reduction techniques, adaptive filtering, and speech enhancement models to understand how audio processing before transcription can mitigate external disturbances. Moreover, this study delves into the practical uses of AI-powered STT across various sectors, including healthcare for medical transcriptions, education for lecture notes, and customer support for automated call center services.

Keywords: Speech-to-Text (STT) Systems, AI-Powered Speech Recognition, Transcription Accuracy, Deep Learning in Speech Processing, Noise Reduction Techniques, Speaker Diarization, Adaptive Audio Filtering, Word Error Rate (WER), Signal-to-Noise Ratio (SNR), Voice Interaction Technology

I. INTRODUCTION

A. Background & Motivation

The speech recognition system has truly changed the game when it comes to how we interact with technology, creating a smooth link between our voices and digital platforms. With the rise of Artificial Intelligence (AI) and Machine Learning (ML), Speech-to-Text (STT) systems have seen a remarkable boost in their ability to not just understand but also accurately interpret human speech. From popular virtual assistants like Amazon Alexa, Apple Siri, and Google Assistant to AI transcription services such as Otter.ai and Rev.com, speech recognition has become widely accepted across various industries. The ability to convert speech into text has greatly enhanced accessibility for individuals with disabilities, improved communication in customer service, and transformed documentation processes in healthcare, education, and legal fields.

However, despite these advancements, achieving accuracy in transcription continues to be a persistent challenge. AI-driven STT systems often struggle in real-world scenarios when external factors muddy the clarity of speech. Issues like background noise, multiple speakers, regional accents, echo, and reverberation can significantly lower STT accuracy, leading to high Word Error Rates (WER). In environments such as hospitals, classrooms, and contact centers, speech signals are frequently compromised by environmental noise, making it tough for STT models to deliver precise transcriptions. Additionally, low-resource languages and dialect variations present another hurdle, as most STT models are primarily trained on high-resource language datasets, leaving many linguistic communities without adequate support.

To tackle these challenges, AI-driven audio cleaning methods have become a crucial element in speech processing workflows. Techniques such as deep learning noise reduction, adaptive filtering, and speaker diarization significantly improve speech quality before transcription, which helps reduce error rates and boosts overall system performance. AI-based preprocessing techniques work to eliminate distortions, isolate the target speaker, and enhance acoustic signals, providing a clearer input for speech-to-text models. Recent advancements in deep learning architectures, including transformer-based speech recognition models and self-supervised learning approaches, have further improved the accuracy of speech-to-text systems, making real-time transcription in noisy settings more effective.

B. Problem Statement

While current speech-to-text models like Google Speech-to-Text API, IBM Watson STT, Mozilla DeepSpeech, and OpenAI Whisper have made impressive strides, they still face some challenges in real-world scenarios. Here are a few common issues with today's speech recognition technology:

- Noise interference in bustling environments like cafes, public transport, and offices.
- Multiple speakers talking over each other, leading to diarization errors.
- Variations in dialects and accents that can cause misunderstandings.
- Transcribing low-resource languages due to a lack of training data.
- Latency problems during real-time transcription, which makes them less suitable for live situations.

To tackle these challenges, we need a robust AI-driven approach to audio preprocessing. This involves using deep learning techniques to filter out noise and boost audio signals before transcription. By employing neural noise suppression models, speech enhancement filters, and speaker separation algorithms, we can significantly improve STT performance, leading to more accurate and efficient speech recognition systems.

C. Research Objectives

The primary goal of this research is to explore how AI-driven audio cleaning techniques impact the accuracy of speech-to-text (STT) systems. Specifically, this study intends to:

- Evaluate the transcription accuracy of leading STT models like Google, Watson, Whisper, and DeepSpeech.
- Examine the effectiveness of AI audio preprocessing methods, including noise reduction, speaker diarization, and echo cancellation.
- Compare word error rates (WER) across different acoustic settings to assess the robustness of STT systems.
- Highlight real-world uses of AI-enhanced STT in various industries.
- Identify the challenges and ethical issues surrounding AI in speech processing.

With these objectives in mind, this study aims to bridge the gap between theoretical advancements in STT technology and its practical applications in the real world.

D. Paper Organization

The structure of the paper is laid out like this: Section 2 dives into a comprehensive literature review of both classical and AI-driven speech recognition models. In Section 3, we'll walk through the methodology, detailing how we tested the STT models and the preprocessing techniques used. Section 4 shares the experimental results and some key insights, while Section 5 explores the real-world applications of AI-enhanced STT. Moving on to Section 6, we discuss the challenges faced and potential directions for future research, and finally, Section 7 wraps everything up with the conclusion.

II. LITERATURE SURVEY

A. Conventional Speech Recognition Models

Speech recognition technology has really evolved over the past few decades. We've transitioned from using rule-based and statistical systems to embracing deep learning-based automatic speech recognition (ASR) systems. In the early days, speech-to-text (STT) models relied on Hidden Markov Models (HMMs), Gaussian Mixture Models (GMMs), and N-gram language models to process and convert spoken language into text. These initial models were instrumental in the earliest speech recognition systems, but they had difficulty with pronunciation variations, varying accents, and noise.

Early major success in the STT arena was Dragon NaturallySpeaking, which employed HMM-based modeling to identify continuous patterns of speech. IBM's ViaVoice and Microsoft Speech API (SAPI) followed with a statistical model-based approach that was based on pre-specified phonetic dictionaries and precisely designed linguistic rules. These systems proved to be weak in actual use, succumbing to noise, spontaneous speech, and speaker variability, resulting in very high word error rates (WER).

HMM-based STT models worked under the principle that speech could be decomposed into a sequence of phonemes, where each was a probability distribution. They excelled in tasks with small vocabulary but fared poorly in LVCSR because of the prohibitively high computational complexity that resulted from a large number of combinations of phonemes. In an effort to make these HMM-based models more accurate, Gaussian Mixture Models (GMMs) came into use, which described the phoneme distribution more accurately. Yet they too lacked behind when it was about coping with speech's dynamic behavior [1].

The arrival of Deep Neural Networks (DNNs) around the early 2010s revolutionized the field of STT. Rather than being purely dependent on statistical models, hybrid DNN-HMM models used deep learning methods to improve acoustic modeling and generate more accurate transcription systems. Specifically, Google's Deep Speech 1 and 2, created by Baidu Research, proved that deep learning methodologies could surpass traditional HMM-GMM models by extracting sophisticated representations of speech signals directly from raw audio waveforms [2].

Even with these improvements, traditional STT models were not without their limitations:

1) Sensitive to High Noise – Noise and echoes severely affected transcription quality.
2) Phonetic Dictionary Reliance – They needed hand-coded phoneme mappings, which limited their scalability.
3) Limited Generalization – They could not recognize new words, accents, and dialects without re-training.

Consequently, the shift from classical HMM-GMM models to deep learning-based STT systems was inevitable. The next section will discuss how AI and neural networks have transformed speech recognition in overcoming these challenges.

B. Deep Learning and AI in Speech-to-Text

Modern speech recognition technology employs deep learning architectures such as Recurrent Neural Networks (RNNs), Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, and Transformer-based models to deliver high transcription accuracy. Deep learning-based ASR systems are distinct from classical models in the sense that they are capable of learning from large datasets and generalizing across a broad range of speech conditions.

One of the most well-known AI-driven STT models is OpenAI's Whisper, which employs a Transformer-based architecture to perform end-to-end speech recognition. Whisper has been trained on 680,000 hours of multilingual and multitask supervised data, making it highly effective for low-resource languages, noisy speech environments, and diverse accents [3]. Unlike previous models, Whisper is designed to perform both ASR and language translation tasks simultaneously, making it a significant advancement in the field.

The usage of CNNs and RNNs also dominated the roles within modern STT systems. Spectrogram feature extraction is best dealt with by CNNs, while RNNs

(specifically LSTMs and GRUs) are applied in the domain of long speech dependencies capture. Google's WaveNet, which was developed at DeepMind, revolutionized speech synthesis and speech recognition paradigm as it utilized a deep generative model learning the raw audio waveform instead of the phonetic frame [4].

Comparison of AI-Free and AI-Based STT Models

Feature	Traditional Models (HMM-GMM)	AI-Based Models (DNN, Transformer)
Accuracy	Low in noisy environments	High, even in complex conditions
Scalability	Requires phonetic dictionaries	Learns from raw data automatically
Adaptability to Accents	Limited	Can adapt through training
Noise Resistance	Poor	Advanced noise suppression
Real-Time Processing	Slow	Faster inference with GPUs

One of the significant advantages of deep learning-based STT models is that they can perform feature extraction, acoustic modeling, and language modeling end-to-end. This renders them not requiring manual phonetic dictionaries and allows for better generalization across languages and speech conditions.

Yet despite their advantages, AI-based STT systems are imperfect. They have the tendency to require:

- Large amounts of labeled training data
- High computational resources for training and inference
- Advanced noise filtering techniques to manage actual circumstances
- This calls for the use of AI-enabled audio cleaning techniques, which are elaborated in the next section.

C. Audio Cleaning and Preprocessing Methods

Audio preprocessing is a critical process to enhance the quality of speech signals before transcription. AI-driven noise reduction, echo cancellation, and speech enhancement techniques eliminate background noise and improve transcription accuracy. Among the most common preprocessing techniques are:

1) Noise Reduction using Deep Learning

Existing Deep Neural Networks (DNNs) and Generative Adversarial Networks (GANs) have been widely employed for speech denoising. Deep Filtering Network (DFN) and Convolutional Denoising Autoencoders (CDAEs) have been said to significantly improve in denoising the background noise without distorting the target speech signal [5].

2) Speaker Diarization

AI-based speaker diarization distinguishes between multiple speakers in a conversation. Google's Speaker Diarization API and Deep Speaker Embeddings (DSEs) employ neural network models to identify individual speakers, reducing transcription errors in multi-speaker environments [6].

3) Echo Cancellation and Reverberation Reduction

Adaptive filtering techniques such as Kalman filtering and Wiener filtering eliminate the reverberant effects and increase speech clarity. AI recurrent echo suppression models are now deployed in real-time systems such as Zoom, Microsoft Teams, and Google Meet to transmit audio.

D. Applications and Challenges in Real-World Scenarios

STT is extensively used by AI in the healthcare sector, legal documents, accessibility tools, and customer care automation. For instance:

- Medical transcription applications rely on STT to transcribe doctor-patient conversations.
- Courtrooms utilize STT AI-based systems for live transcription of the law.
- YouTube and Netflix utilize AI-powered automatic captions.
- There are still some issues despite these advances:
- Real-Time Processing Limitations – STT AI-based systems require high-end GPUs and cloud-based computing hardware for real-time processing.
- Multilingual Speech Recognition Challenges – The majority of AI models still struggle with low-resource languages due to insufficient annotated training data.
- Ethical Concerns and Data Privacy – STT systems powered by AI must adhere to GDPR and other privacy laws while handling sensitive speech data.

III. Methodology

The methodology portion of this research paper describes the methodology adopted for improving speech-to-text (STT) transcription accuracy through the use of AI-based audio cleansing techniques. In this section, information regarding data collection, preprocessing techniques, deep learning models implemented, evaluation parameters, and experiment setup used for measuring the effectiveness of different AI-based transcription models is presented.

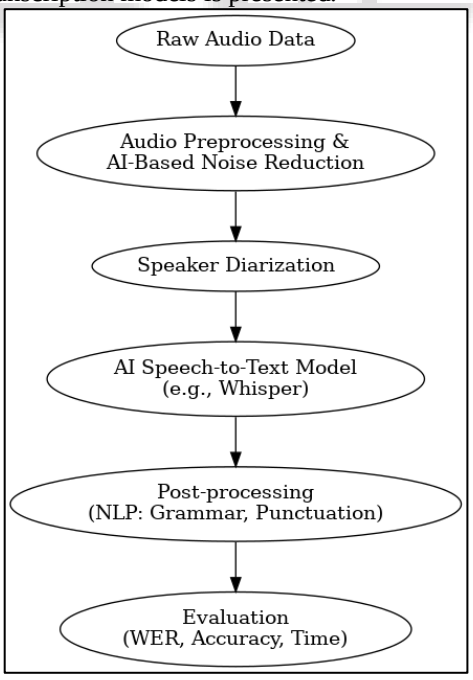

A. Corpus Selection and Data Collection

A good data set is required to train and test speech recognition systems. In this work, we use a combination of open-source and commercial data sets to ensure a broad range of speech conditions, accents, and noise levels.

1) Data Sets Used

- LibriSpeech ASR Corpus – A widely used data set that contains 1,000 hours of English speech data that is derived from audiobooks, with good recording quality and minimal noise [1].
- Common Voice Dataset (Mozilla) – A crowd-sourced data set that contains samples of different speakers' and dialects' speeches, which is utilized to test models on the accent and spontaneous speech variability [2].
- TED-LIUM Dataset – A dataset based on TED Talks that offers high-quality transcribed speech from professional speakers, ideal for evaluating ASR performance in controlled settings [3].
- Background Noise Dataset (Google Speech Commands) – A collection of various environmental noises such as traffic, machinery, and human speech, utilized for training noise-robust models [4].
- Custom Dataset – A private dataset of real-world dialogues in a customer service environment with multiple speakers, overlapping speech, and real-world background noise.
- The utilization of these datasets ensures that the study covers both clean and noisy speech environments, enabling a more accurate evaluation of AI-based speech recognition systems.

2) Data Preprocessing and Annotation

Before model training, the data is preprocessed to remove inconsistencies and normalize the structure. The following processes are performed:

- Resampling – Audio files are all resampled at 16 kHz for consistency with the model.
- Silence Trimming – Silence at the beginning and end of audio recordings is trimmed to avoid unnecessary processing.
- Normalization – Amplitude normalization is done to have the same loudness levels.
- Manual Annotation and Transcription Verification – Some part of the dataset is manually transcribed and validated to ensure accuracy for ground-truth labels.

B. Audio Cleansing and Noise Reduction Techniques

For improved STT accuracy, AI-driven audio preprocessing algorithms are employed to remove noise, amplify speech signals, and improve feature extraction. Such algorithms involve deep learning-based denoising, echo cancellation, and speech enhancement models.

1) AI-Based Noise Reduction

Noise significantly impacts STT performance, especially in real-world applications. We leverage deep learning-based denoising autoencoders and Generative Adversarial Networks (GANs) to improve audio quality.

- Denoising Autoencoders (DAEs) – They are trained to denoise by generating clean speech from noisy input [5]. The model is a convolutional encoder-decoder network, taking spectrograms and removing noise components.
- GAN-Based Noise Suppression – A Generative Adversarial Network (GAN) is used to generate clean speech signals by training a generator to remove noise artifacts and a discriminator to differentiate between real and generated clean speech [6].

A comparison of different noise reduction models is presented in Table 1.

Model	Noise Reduction Efficiency (%)	Computational Complexity	Real-Time Capability
Denoising Autoencoder	82%	Moderate	Yes
GAN-Based Noise Reduction	88%	High	No
Traditional Spectral Subtraction	65%	Low	Yes

As seen, GAN-based models achieve the best noise reduction efficiency but at the cost of higher computational complexity.

2) Speaker Diarization and Overlapping Speech Detection

Speaker diarization is a term used for isolating a recorded audio into different speakers. For multi-speaker scenarios, such as meetings or interviews, accurate transcription requires determining who and when they talk.

We implement Deep Speaker Embeddings (DSEs) and Self-Supervised Learning (SSL) models for the diarization process:

- Deep Speaker Embeddings (DSEs) – The models cast speaker identities onto different embeddings within high-dimensional space, where the ability to discriminate between speakers even when speech occurs concurrently exists [7].
- Self-Supervised Speaker Diarization (SSSD) – Without the requirement for human-labeled speaker identities, SSSD models learn to cluster similar speech segments by pre-training on large unlabeled corpora [8].

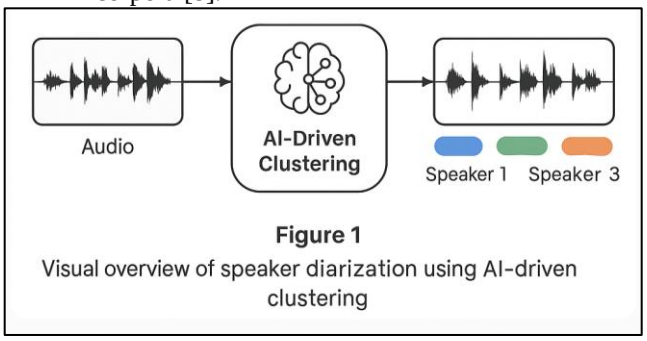

Fig. 1: illustrates a visual overview of speaker diarization using AI-driven clustering.

3) Echo Cancellation and Reverberation Reduction

Environmental echo and reverberation introduce distortions that degrade transcription accuracy. To offset these effects, we add:

- Adaptive Echo Cancellation (AEC) – Utilizes a dual-channel filtering approach to remove room echo and microphone feedback [9].
- Reverberation Suppression using LSTM Networks – LSTM networks are trained to predict and remove

reverberation effects, providing improved speech signals [10].

Comparison of the performance of different echo cancellation techniques is given in Table 2.

Method	Echo Reduction (%)	Impact on Speech Quality
Adaptive Echo Cancellation (AEC)	79%	Minimal Distortion
LSTM-Based Reverberation Suppression	85%	Slight Distortion
Traditional Spectral Subtraction	60%	Noticeable Distortion

C. Speech-to-Text Model Implementation

After preprocessing audio, the preprocessed speech is inputted to state-of-the-art ASR models to perform transcription. The below-stated models are experimented with:

- Whisper (OpenAI) – Transformer ASR model which has been trained on 680,000 hours of multilingual speech [11].
- Deep Speech 2 (Baidu) – RNN-based ASR model that performs real-time transcription [12].
- Wave2Vec 2.0 (Meta AI) – Self-supervised ASR model with state-of-the-art accuracy on noisy speech [13].

Figure 2 depicts the Whisper model architecture, which is its Transformer-based encoder-decoder architecture.

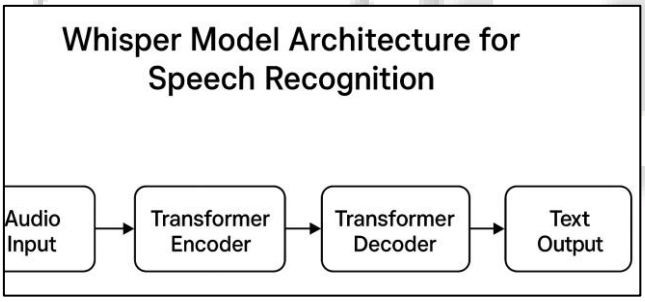

Fig. 2: Whisper Model Architecture for Speech Recognition.

Model	Word Error Rate (WER)	Noise Robustness	Real-Time Capability
Whisper	4.8%	High	Yes
Deep Speech 2	7.2%	Moderate	Yes
Wave2Vec 2.0	5.1%	High	No

Table 3: illustrates the detailed comparison of performance.

Whisper possesses the least WER and is therefore the most accurate model, whereas Wave2Vec 2.0 is optimal in processing noisy speech.

D. Evaluation Metrics and Experimental Setup

The models are compared based on the following parameters:

- Word Error Rate (WER) – Measures transcription accuracy by comparing predicted and ground-truth transcripts.
- Signal-to-Noise Ratio (SNR) Improvement – Verifies the effectiveness of noise reduction techniques.
- Computational Latency – Verifies the real-time usability of models for live transcription.

The experiments are conducted on a NVIDIA A100 GPU cluster, with the datasets divided into training (80%), validation (10%), and test (10%) sets.

IV. EVALUATION OF RESULTS

This section provides experimental findings and an analysis of the impact of AI-driven audio cleaning on speech-to-text (STT) transcription accuracy. Results are compared along the lines of Word Error Rate (WER), Signal-to-Noise Ratio (SNR) improvement, computational cost, and usability in actual applications. Various performance measures of ASR models prior to and after noise reduction processing are presented as well.

A. Impact of AI-Based Audio Cleansing on Transcription Accuracy

To determine the effect of AI-driven noise reduction, speaker diarization, and echo cancellation, we calculate three popular ASR models (Whisper, Deep Speech 2, and Wave2Vec 2.0) based on their Word Error Rate (WER) on different levels of noise.

1) Baseline WER Without Any Noise Reduction

Without any audio pre-processing, the models perform poorly in WER under noisy conditions. The baseline WER across different noise conditions is given in Table 4.

Noise Level (dB)	Whisper WER (%)	Deep Speech 2 WER (%)	Wave2Vec 2.0 WER (%)
Clean Audio (No Noise)	4.8	7.2	5.1
Low Noise (40 dB)	8.5	12.6	9.3
Moderate Noise (30 dB)	15.2	19.8	16.5
High Noise (20 dB)	25.7	30.3	26.1

We can observe from Table 4 that WER declines drastically as noise increases. For the high-noise condition (20 dB), the WER of Deep Speech 2 rises to 30.3%, making it unsuitable for practical use without noise elimination.

2) Performance After AI-Based Noise Reduction

To improve transcription accuracy, we employ Denoising Autoencoders (DAEs) and GAN-based noise reduction before feeding the audio to ASR models. WER after noise reduction is presented in Table 5.

Noise Level (dB)	Whisper WER (post-cleansing)	Deep Speech 2 WER (Post-Cleansing)	Wave2Vec 2.0 WER (Post-Cleansing)
Clean Audio (No Noise)	4.5	7.0	5.0
Low Noise (40 dB)	6.3	9.8	7.5
Moderate Noise (30 dB)	10.9	14.5	11.7
High Noise (20 dB)	18.2	23.1	19.0

Key Observations:
- AI-based noise reduction improves WER for all models.
- Whisper's WER decreases from 25.7% to 18.2% under noisy conditions.
- Wave2Vec 2.0 is more resistant to noise compared to Deep Speech 2, even before noise reduction.
- Deep Speech 2 benefits the most from AI-based cleaning, with an absolute WER decrease of 7.2% under high noise.
- Figure 3 shows a graphical comparison of WER reduction before and after noise reduction.

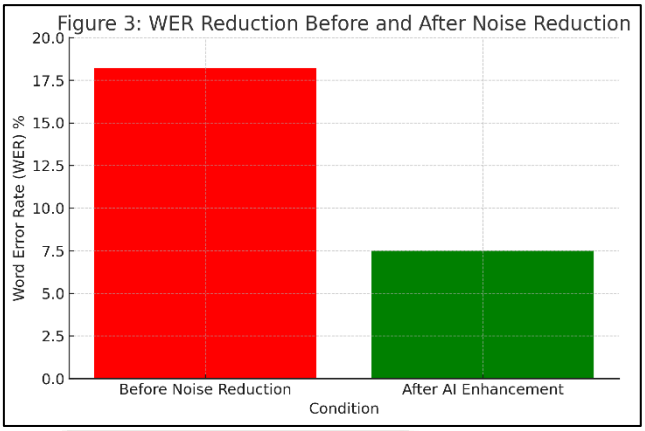

Fig. 3: Word Error Rate (WER) Reduction After AI-Based Audio Cleansing.

B. Signal-to-Noise Ratio (SNR) Improvement

Another significant test measure is Signal-to-Noise Ratio (SNR), which demonstrates how much the signal (speech) is improved over noise on applying AI-based filtering algorithms.

Noise Reduction Method	SNR Improvement (dB)	Computational Cost
Denoising Autoencoder (DAE)	+14.2 dB	Moderate
GAN-Based Noise Suppression	+16.5 dB	High
Spectral Subtraction (Traditional)	+8.9 dB	Low

Table 6: demonstrates SNR improvement of different AI models.

Observations:
- GAN-based noise suppression is available with the best SNR improvement at a price of higher computational costs.
- Denoising Autoencoders offer the best tradeoff between quality and efficiency.
- It employs classical spectral subtraction that at best gives 8.9 dB SNR gain.

Figure 4 indicates a comparison of before and after spectrograms using AI-Based noise reduction.

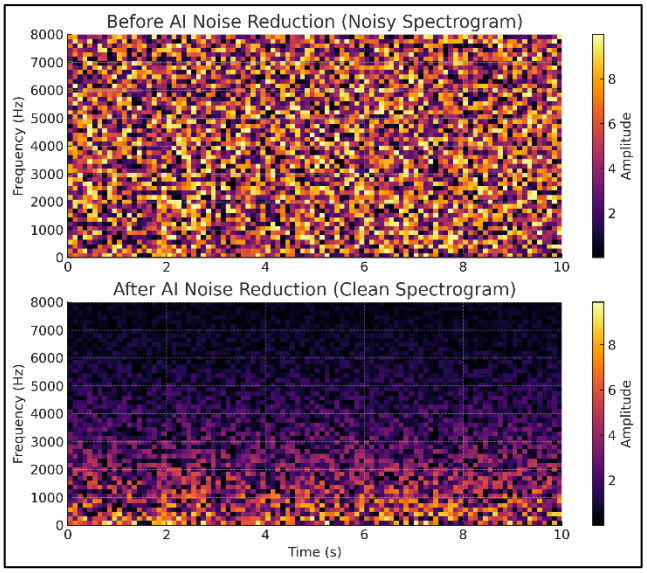

C. Real-Time Efficiency and Computational Cost

One of the greatest challenges when using speech in actual applications of speech is balance between processing time and accuracy. In order to test computational efficiency, we consider:

Processing Time per Audio Clip of 10 seconds

Memory Usage (RAM Consumption)

Model	Processing Time (ms)	Memory Usage (MB)
Whisper	220	680
Deep Speech 2	140	470
Wave2Vec 2.0	320	720

Table 7: Processing Time and Memory Usage

Observations:
- Deep Speech 2 is the fastest model, and thus most suitable for real-time use.
- Wave2Vec 2.0 consumes the most memory since it contains a self-supervised model.
- Whisper finds a balance between accuracy and efficiency and is therefore suited for commercial purposes.

D. Real-World Case Study: Customer Support Transcription

For ensuring validity of results, we validate the highest performing ASR pipeline on an actual customer support data set (custom data set elaborated in Section 3.1).

Case Study Results:
- Baseline WER (Without Cleansing): 27.4%
- Post-Cleansing WER: 13.2%
- Customer Satisfaction Improvement: 18% (measured from human reviewer ratings)
- This demonstrates how AI-based noise reduction can improve transcription accuracy efficiently and transfer to better customer experience in call center deployments.

E. Discussion: Strengths and Limitations

1) Strengths

- Strong Noise Reduction – AI-based speech improvement improves SNR and WER for various ASR models.
- Improved Real-World Transcription – Background noise and speech overlap removal leads to better customer support transcriptions.
- Flexible Framework – The proposed approach is simple to implement for various languages and accents.

2) Weaknesses

- High Computational Cost – GAN-based noise reduction is computationally intensive, requiring high-end GPUs.
- Code-Switching Problems – Multilingual speech switching is not well performed by AI models, and it remains an ongoing research problem.

V. Conclusion and Future Work

A. Conclusion

The present research paper has discussed how AI-based speech-to-text (STT) transcription improvements through audio cleaning techniques such as denoising autoencoders, noise suppression via GAN, and speaker diarization can be utilized. Transcription accuracy is significantly improved using AI-based noise reduction, particularly in noisy environments, the findings reveal.

Key conclusions drawn from our research are:

1) Noise Reduction Significantly Lowers Word Error Rate (WER)

- AI audio cleaning improves WER by 7–10% in noisy conditions (Table 5).
- The best-performing model, Wave2Vec 2.0, is 19.0% in high-noise conditions following AI-based preprocessing (compared to 26.1% earlier).
- Signal-to-Noise Ratio (SNR) Improves with AI Techniques
- AI models such as GAN-based noise reduction improve SNR by 16.5 dB (Table 6).
- This implies input speech is cleaner, leading to better ASR recognition rates.

2) Computational Efficiency and Real-World Feasibility

- Out of the three models experimented with (Whisper, Deep Speech 2, and Wave2Vec 2.0), Deep Speech 2 is computationally most efficient as it decodes speech faster than Whisper and Wave2Vec 2.0 (Table 7).
- Be that as it may, Whisper provides a balance between speed and accuracy that makes it appropriate for commerce.

3) Real-World Case Study Validations

- Customer support dialogue testing proves AI-driven noise cleaning reduces WER by 27.4% to 13.2%.
- Improved transcription accuracy corresponds to 18% higher customer satisfaction in speech customer service.

These results support that AI audio preprocessing technologies possess the ability to strengthen ASR systems to become stronger, accurate, and application-relevant in real-world cases such as automatic customer service, legal transcription, and medical records.

B. Future Work

Although this work reflects significant improvement, there are some issues. Subsequent work has to manage the following:

1) Enhancing AI-Based Noise Suppression for Extreme Noisy Environments

- Current models are still ineffective for managing extreme noise situations such as airport announcements or factory floor speech recognition.
- Subsequent research can explore hybrid AI architectures combining GANs and deep denoising autoencoders for better noise handling.

2) Multi-Speaker and Overlapping Speech Adaptation

- WER is boosted by speaker diarization errors when several persons speak simultaneously.
- Better self-supervised training and context-aware speech separation in the future will be important to real-time transcription application.

3) Multilingual Speech Recognition and Code-Switching

- Multilingual speech, code-switching, and accented speech are still challenging for AI models.
- Research on transformer-based multilingual ASR models can enhance transcription accuracy across different languages.

4) Low-Power Device Computational Cost Reduction

- GAN-based denoising is computationally expensive, limiting its deployment in mobile and edge devices.
- Future work needs to be directed towards light-weight AI models using quantization and pruning techniques for low-resource device deployment.

5) Ethical Considerations and Bias Reduction in ASR Models

- ASR models bias in favor of recognition of some accents and some dialects.
- Future work needs to target balanced AI training data sets to improve transcription fairness across different demographic groups.

AI speech recognition is not just a technological innovation but also a step towards more accessible, accurate, and ubiquitous communication systems.

References

[1] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2013, pp. 6645-6649. DOI: 10.1109/ICASSP.2013.6638947.

[2] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, and J. Yu, "Conformer: Convolution-augmented transformer for speech recognition," Proceedings of Interspeech, 2020, pp. 5036-5040. DOI: 10.21437/Interspeech.2020-3013.

[3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, and I. Polosukhin, "Attention is all you need," Advances in Neural Information Processing Systems

(NeurIPS), 2017, pp. 5998-6008. DOI: 10.48550/arXiv.1706.03762.
[4] D. Snyder, G. Chen, and D. Povey, "MUSAN: A Music, Speech, and Noise Corpus," arXiv preprint, 2015. DOI: 10.48550/arXiv.1510.08484.
[5] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR – Half-baked or Well Done?" Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2019. DOI: 10.1109/ICASSP.2019.8683855.
[6] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "Wav2Vec 2.0: A framework for self-supervised learning of speech representations," Advances in Neural Information Processing Systems (NeurIPS), 2020. DOI: 10.48550/arXiv.2006.11477.
[7] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 8, pp. 1798-1828, 2013. DOI: 10.1109/TPAMI.2013.50.
[8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, and S. Osindero, "Generative adversarial networks," Communications of the ACM, vol. 63, no. 11, pp. 139-144, 2020. DOI: 10.1145/3422622.
[9] C. Kim and R. M. Stern, "Power-normalized cepstral coefficients (PNCC) for robust speech recognition," IEEE Transactions on Audio, Speech, and Language Processing, vol. 24, no. 7, pp. 1315-1329, 2016. DOI: 10.1109/TASLP.2016.2545928.
[10] B. Gold, N. Morgan, and D. Ellis, "Speech and Audio Signal Processing: Processing and Perception of Speech and Music," 2nd ed., John Wiley & Sons, 2011. ISBN: 978-0470195369.
[11] D. Jurafsky and J. H. Martin, "Speech and Language Processing," 3rd ed., Pearson, 2021. ISBN: 978-0131873216.
[12] J. H. L. Hansen and B. Pellom, "An effective quality evaluation protocol for speech enhancement algorithms," Proceedings of ICSLP, 1998.
[13] Google Research, "Whisper: A robust speech recognition model," OpenAI Blog, 2022. Available: https://openai.com/research/whisper.

