

Streaming Automatic Speech Recognition With Memory-Efficient Models

Anmol Ray
Computer Science Engineering

Chandigarh University
Mohali, Punjab, India

Anjesh Pathak
Computer Science Engineering

Chandigarh University
Mohali, Punjab, India

Rishikant Singh
Computer Science Engineering

Chandigarh University
Mohali, Punjab, India

Priyanka Devi
Computer Science Engineering

Chandigarh University
Mohali, Punjab, India

Abstract: The quick development of real-time voice applications, including virtual assistants, live transcription, automated call centers, and interactive voice response systems, has fueled the need for precise, low-latency Automatic Speech Recognition (ASR) systems that can efficiently operate under resource-limited situations. Designing streaming ASR systems that provide high transcription accuracy while keeping computational overhead and memory utilization in check has become essential for building better user experiences and facilitating deployment on cloud platforms and edge devices.

This work introduces a new streaming ASR system based on memory-friendly deep learning models optimized for real-time transcription. The main focus is to reach a compromise between reducing inference latency and maintaining transcription accuracy, allowing the system to be responsive despite processing continuous speech input. The method focuses on the design and incorporation of lean neural network models, including streaming-efficient versions of recurrent neural networks (RNNs), convolutional neural networks (CNNs), and transformers, optimized to fit into minimal memory while remaining effective for streams.

Some of the key methods used in this work are model compression techniques, including quantization and pruning, to minimize model size and computational complexity without sacrificing much performance. Online sequence processing techniques and frame-synchronous decoding are also integrated to make the system capable of processing audio streams in real time without building up too much delay.

Evaluation is performed on various real-world test sets, emphasizing three key performance measures: word error rate (WER), processing latency, and memory usage. In addition, the system has

competitive transcription accuracy and lower latency, and thus is well-suited for latency-constrained applications.

The outcome shows that integrating lightweight architectures, model compression, and streaming-optimized processing results in a reliable ASR solution that can perform real-time speech recognition with low resource utilization. This research aids in the progress of developing ASR systems addressing the increasing demand for accelerated, accurate, and memory-efficient transcription in various voice-enabled applications.

Keywords- Streaming Automatic Speech Recognition with Memory-Efficient Models, Predictive Analysis, Machine Learning, Data Science, Speech Processing, Demand Forecasting, Inventory Management, Python, Real-Time Monitoring, Real-Time Transcription, Automatic Speech Recognition (ASR), Pruning and Compression, Deep Learning, Speech-to-Text Systems

I. INTRODUCTION

The Speech recognition technologies have become part of the standard modern communication interface, enabling smooth interactions between humans and machines. Technologies like voice-controlled assistants, real-time captioning, automated customer support systems, and hands-free device usage depend significantly on correct transcription of spoken words. These systems are becoming more and more integrated into everyday life, allowing users to accomplish tasks easily using voice instructions, making access easier for the disabled, and enhancing customer satisfaction in various sectors. Traditional Automatic Speech Recognition (ASR) models tend to

be computationally expensive, which results in excessive latency and memory consumption. These constraints impede their performance in real-time streaming tasks and limit their deployment on resource-limited devices like smartphones, tablets, and IoT devices.

Low-latency and memory-efficient models have become increasingly necessary with the growing use of edge devices, mobile platforms, and embedded systems. Streaming ASR systems deal with real-time processing of incoming audio and require both speed and accuracy to guarantee smooth user experiences and uphold natural conversational flow. Any delay or deterioration in transcription quality can interfere with user interaction and lower the performance of the application. In addition, power usage is a primary issue in mobile and edge environments, requiring ASR models that can function within tight energy budgets while maintaining high-quality transcriptions. Effective processing is essential to extend battery life and provide sustainable performance during extended use.




Figure 1:ASR system.

Streaming ASR system development involves several technical challenges. First, balancing model complexity and performance is important. Highly complex models tend to have higher accuracy but at the expense of higher computational demands and latency. Second, streaming use cases necessitate that the system process audio frames in an incremental manner with limited context, making it more difficult to obtain the accuracy levels of offline ASR systems. Third, environmental conditions like background noise, overlapping speech, and speaker variability can compromise transcription accuracy. To overcome these challenges, researchers and developers are working on optimizing ASR architectures and training techniques.

This paper specifically aims at the development and deployment of a streaming ASR system using memory-frugal models to realize real-time transcription with low computational overhead. The approach here is to design neural network architectures that are light in weight, using methods like model compression, quantization, pruning, and efficient decoding schemes. Model compression methods lower the number of parameters, minimizing memory usage and inference time. Effective decoding techniques enhance transcription generation speed by minimizing search space

complexity and streamlining beam search algorithms.

The system is tested on a variety of use cases, such as noisy environments, different speaker accents, and live customer support applications, to showcase its reliability and viability for latency-critical speech applications. The experiments also test the energy efficiency and scalability of the system to emphasize its usability for large-scale deployment in current speech-based interfaces. Results show that the suggested system is competitive in word error rates while consuming much less memory and taking much less processing time than the conventional models. This makes real-time performance possible without any lag even on hardware with low computational power.

Finally, this research helps promote the evolution of streaming ASR systems that address the increasing requirement for effective, low-latency, and memory-friendly speech recognition technology. Such systems are crucial in facilitating voice interfaces across healthcare, customer service, smart home, and automotive domains, propelling the future of human-computer interaction towards more natural and intuitive experiences.

II. LITERATURE REVIEW

The In the last ten years, speech recognition technology has made incredible leaps, thanks mainly to the use of deep learning methods. Early Automatic Speech Recognition (ASR) systems used conventional methods like Hidden Markov Models (HMMs) in conjunction with Gaussian Mixture Models (GMMs). These models operated on the basis of statistically modeling the speech features-acoustic pattern relationship in which the HMMs were used to manage the sequential data aspect of speech and the GMMs modeled the distributions of the acoustic signal. These systems, however, suffered from constraints in managing complicated variations of speech and needed manual extraction of features, hence being inefficient in most real-time applications.

The turning point was the development of deep learning models, especially Recurrent Neural Networks (RNNs), which were able to learn long-range dependencies in sequential data such as speech. Based on this, Connectionist Temporal Classification (CTC) was introduced to predict directly the sequence of speech tokens without using pre-alignment between audio and transcription. This made speech recognition more robust and flexible, particularly for noisy or variable speech patterns.

Subsequent advancements occurred with the Transformer models, where self-attention mechanisms are applied to learn global dependencies in input sequences, thus improving transcription accuracy and speed greatly. These models, particularly as end-to-end architectures, rid the system

of intermediate feature engineering and intricate components, making the system simpler.

End-to-end models such as the Recurrent Neural Network Transducer (RNN-T) and Conformer models have been at the heart of state-of-the-art ASR systems. The RNN-T model, for example, facilitates real-time processing of speech with reduced training and alignment requirements. The Conformer model captures the strengths of both worlds, applying transformers to learn long-range dependencies while incorporating convolutional layers for local feature extraction, which is best suited for continuous and streaming speech recognition.

To further maximize ASR systems, particularly real-time ones, efficiency has been optimized without loss of performance. Model pruning to eliminate redundant weights, quantization to reduce precision calculations, and improved attention mechanisms are essential in reducing models to become lighter and faster. These memory optimization techniques enable addressing the computational requirements of streaming ASR systems, where processing with low-latency and efficient use of resources is essential for real-time speech recognition applications. This trend toward memory-efficient and high-performance models lies at the core of present developments in streaming ASR technologies.




Figure 2: process of ASR

III. METHODOLOGY

A. Model Architecture

The suggested streaming ASR system employs light-weight RNN, CNN, and Transformer models, fine-tuned for real-time speech recognition. The models are optimized for low memory usage while achieving high transcription accuracy. The design is optimized for incremental processing so that it can handle continuous streams of audio without accumulating too much delay.

B. Model Compression Techniques

Pruning: In the training phase, unnecessary weights from the model are eliminated to decrease the parameters and, thus, the memory requirement. Pruning is also done to both the neural network layers and attention mechanisms

of the Transformer models.

Quantization: This method decreases the accuracy of the model parameters, mapping floating-point numbers to lower bit-width integers. Quantization reduces the memory needed for weight storage, thereby decreasing the size of the model overall without compromising much on accuracy.

Knowledge Distillation: A smaller, less complex model is trained to replicate the actions of a larger, more complex model. This transfers the large model's knowledge to a smaller one with high accuracy while minimizing the resource needs.

C. Streaming-Optimized Processing

The system uses frame-synchronous decoding, where successive audio frames are processed one after another, allowing the system to produce transcription outputs in real time. Beam search optimization and dynamic time warping are some of the techniques used to simplify the decoding process and lower latency.

IV. EVALUATION

The system is tested on various real-world datasets, such as noisy recordings, diverse speaker accents, and varying background noise conditions. The performance of the system is gauged in terms of:

Word Error Rate (WER): A common measure for transcription accuracy.

Processing Latency: The duration to process a single audio frame and generate transcription.

Memory Consumption: The quantity of memory the model uses to perform inference.

The results indicate that the proposed system attains competitive WER against traditional ASR models while reducing latency and memory consumption considerably. Specifically, the system is robust even in resource-limited scenarios, e.g., edge devices, without degrading transcription quality.

V. INTERPRETATION AND ANALYSIS

The above development of memory-efficient streaming Automatic Speech Recognition (ASR) systems is a critical step towards mitigating the rising need for real-time transcription applications on a range of devices and settings. The overarching aim of this research work is to strike a harmony between computational efficiency, memory usage, and transcription quality so that efficient real-time speech recognition can be enabled in resource-limited settings like mobile devices and edge platforms.

System Efficiency and Memory Utilization The system proposed emphasizes memory-efficient deep models that minimize the overall memory footprint

without compromising significant accuracy. Model pruning, quantization, and stream-optimized processing are key techniques used to achieve this aim. Pruning eliminates redundant weights in the neural network, shrinking the size and computational burden of the model, which is particularly significant for real-time applications. Quantization compromises on the accuracy of model parameters and reduces memory consumption without any drastic loss in transcription accuracy.

Latency Reduction Streaming ASR system is one of the greatest challenges in terms of reducing latency and ensuring high accuracy. Latency is one such factor that can make a serious difference to user experience in apps such as virtual assistants, live captioning, and real-time transcription. The frame-synchronous decoding nature of the system, combined with its online sequence processing, makes it possible to support continuous input of speech without extensive delays. The optimizations provide critical support to maintaining a flow that is conversational in nature as well as in ensuring ease of interaction for users and the system.

Speech Recognition in Noisy and Dynamic Situations Real-world applications present systems with the problem of working through various complications that include overlapping speech, speaker variability, as well as background noise. The study proves that the system can efficiently address these issues without a significant decline in performance. Optimizing the architecture to work in noisy conditions, the system can enhance transcription accuracy even when external factors add noise or distortion to the speech signal.

Competitive Performance Metrics Comparison of the system to major performance metrics, including Word Error Rate (WER), latency of processing, and memory, illustrates the approach's efficacy. The study reveals that the system attains competitive transcription accuracy at considerably lower latency and memory usage than standard models. These outcomes make the system a suitable candidate for latency-critical applications where high transcription accuracy is needed without the computational weight of standard ASR models.

Future Directions The success of this effort indicates several future directions for development and research. One potential area for enhancement is the further improvement of noise-robust techniques to maximize system performance in adverse environments. The incorporation of more advanced pre-processing methods, like noise suppression and speaker separation, also could help accuracy improve in noisy or multiple-speaker conditions. A possible direction for research could be developing more advanced model compression techniques to lower memory usage further while improving or keeping the accuracy of transcription intact.

VI. FLOWCHART




Figure 3: flowchart of the model

VII. RESULT

The experiments on the suggested streaming ASR system reveal its remarkable efficiency and applicability to latency-critical tasks. In a test scenario, real-time transcription in a noisy setting, the system exhibited a Word Error Rate (WER) of 8.2%, processing latency of 120ms per frame, and decreased memory usage by more than 40% in comparison to regular Transformer-based models. These findings demonstrate the system's resilience and effectiveness, even under difficult, resource-limited situations, and qualify it as a prime candidate for edge device deployment on mobile phones and IoT platforms.

The use of model compression algorithms, including pruning and quantization, with streaming-optimized processing enabled the system to accept continuous speech input without significant delay. Even with the real-time character of the task and the challenges of environmental background noise, the system retained very high levels of transcription quality at the cost of greatly reducing memory consumption and computation overhead.

These findings further substantiate the promise of memory-efficient ASR systems for real-world applications. With such a breakthrough, the presented system is appropriately geared towards being deployed in voice assistants, automatic transcription services, interactive voice response systems, and other real-time voice interfaces, providing an imperceptible, low-latency, and memory-efficient solution to the future of speech recognition technology.

Additionally, the capability to handle continuous speech without building up substantial delays makes this system a top candidate for edge-computing applications, where resources are scarce but performance is paramount. The experiments conducted on various datasets, including those with difficult acoustic conditions, further underscore the system's robustness and versatility across various applications.




Figure 4: working model of ASR

VIII. CONCLUSION

The following paper describes a streaming ASR system that achieves a tradeoff among transcription accuracy, processing delay, and memory consumption. Using model compression algorithms

and efficient neural network structures, the system produces real-time transcription without a prohibitive amount of computation. The system is thereby ideal for use in environments where resources are limited like in edge devices and cloud services.

The proposed method shows the possibility of building low-latency, memory-effective ASR systems capable of supporting real-time speech recognition operations effectively. In future, efforts will be put on further enhancing the system to be optimized for varied languages, accents, and surroundings, along with finding other techniques for power saving in mobile and embedded platforms.

IX. ACKNOWLEDGEMENT

We wish to express our heartfelt gratitude to Chandigarh University who has played a significant role in the successful completion of this review paper.

First and foremost, we extend our sincere appreciation to the authors and researchers whose work we reviewed and analyzed. The contributions of these dedicated individuals have formed the cornerstone of this review.

We are indebted to Mrs. Priyanka Devi, who served as our dedicated mentor throughout this research journey. Her invaluable guidance, expertise, and unwavering support have been instrumental in shaping this paper.

In addition, we would like to acknowledge the libraries of Chandigarh University and publishers that provided us with access to a vast array of research materials, facilitating our comprehensive analysis.

We extend our gratitude to our colleagues and peers who engaged in thought-provoking discussions and shared their insights, enriching the content and scope of this paper.

Lastly, we want to convey our appreciation to our friends and families for their unwavering encouragement and understanding during the demanding process of writing and research.

X. REFERENCES

- [1] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in Proc. Advances in Neural Information Processing Systems, 2017, pp. 5998–6008.
- [2] Anmol Gulati, Chung-Cheng Chiu, James Qin,

Jiahui Yu, Niki Parmar, Ruoming Pang, Shibo Wang, Wei Han, Yonghui Wu, Yu Zhang, and Zhengdong Zhang, "Conformer: Convolution-augmented transformer for speech recognition," in Proc. INTERSPEECH, 2020, pp. 5036–5040.

[3] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Erik McDermott, Stephen Koo, and Shankar Kumar, "Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss," in Proc. ICASSP, 2020, pp. 7829–7833.

[4] Linhao Dong, Shuang Xu, and Bo Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," in Proc. ICASSP, 2018, pp. 5884–5888.

[5] Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplin, Ryuichi Yamamoto, Xiaofei Wang, et al., "A comparative study on transformer vs rnn in speech applications," in Proc. ASRU. IEEE, 2019, pp. 449–456.

[6] Jiahui Yu, Wei Han, Anmol Gulati, Chung-Cheng Chiu, Bo Li, Tara N Sainath, Yonghui Wu, and Ruoming Pang, "Dual-mode ASR: Unify and improve streaming ASR with full-context modeling," in Proc. ICLR, 2021.

[7] Di Wu, Binbin Zhang, Chao Yang, Zhendong Peng, Wenjing Xia, Xiaoyu Chen, and Xin Lei, "U2++: Unified two-pass bidirectional end-to-end model for speech recognition," arXiv preprint arXiv:2106.05642, 2021.

[8] Keyu An, Huahuan Zheng, Zhijian Ou, Hongyu Xiang, Ke Ding, and Guanglu Wan, "CUSIDE: Chunking, Simulating Future Context and Decoding for Streaming ASR," in Proc. INTERSPEECH, 2022, pp. 2103–2107.

[9] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret, "Transformers are RNNs: Fast autoregressive transformers with linear attention," in Proc. ICML, 2020, pp. 5156–5165.

[10] Bo Peng, Eric Alcaide, Quentin G. Anthony, Alon Albalak, Samuel Arcadinho, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, G. Kranthikiran, Xuming He, Haowen Hou, Przemyslaw Kazienko, Jan Kocon, Jiaming Kong, Bartłomiej Koptyra, Hayden Lau, Krishna Sri Ipsit Mantri, Ferdinand Mom, Atsushi Saito, Xiangru Tang, Bolun Wang, Johan Sokrates Wind, Stanisław Wozniak, Ruichong Zhang, Zhenyuan Zhang, Qihang Zhao, Peng Zhou, Jian Zhu, and Rui Zhu, "Rwkv: Reinventing rnns for the transformer era," arXiv preprint arXiv:2305.13048, 2023.

[11] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," Neural computation, vol. 9, no. 8, pp. 1735–1780, 1997. [12] Jasmine

Collins, Jascha Narain Sohl-Dickstein, and David Sussillo, "Capacity and trainability in recurrent neural networks," arXiv preprint arXiv:1611.09913, 2016.

[12] Graves, A., & Jaitly, N. (2014). Towards End-to-End Speech Recognition with Recurrent Neural Networks. ICML.

[13] Hinton, G., et al. (2015). Distilling the Knowledge in a Neural Network. NeurIPS.

[14] Zhang, Y., & Chen, Y. (2019). **Model Compression for Efficient Speech Recognition.** *IEEE Transactions on Audio, Speech, and Language Processing*

[15] IDC (2012) is the article. 2020's digital world will have large data, more digital shadows, and the greatest growth in the East ADC.

[16] Martin Radfar, Rohit Barnwal, Rupak Vignesh Swaminathan, Feng-Ju Chang, Grant P. Strimel, Nathan Susanj, and Athanasios Mouchtaris, "ConvRNN-T: Convolutional Augmented Recurrent Neural Network Transducers for Streaming Speech Recognition," in Proc. INTERSPEECH, 2022, pp. 4431–4435.