IOCMA 2026 Conference

The 2nd International Online Conference on Mathematics and Applications

10–12 June 2026 | Online

MDPI

Algorithm Design and Mathematical Modeling for Efficient Automatic Speech Recognitionin Low-Resource African Languages

Simanga Comfort Mchunu

Department of Mathematical Sciences · 47483571@mylife.unisa.ac.za · University of South Africa, Pretoria · 2026

INTRODUCTION & AIM

Deploying Automatic Speech Recognition (ASR) for African languages is critically hindered by the incompatibility of large-scale models with mobile hardware. Models like NLLB exceed 1 GB in size and cannot load on typical devices with 2–4 GB RAM, preventing real-time conversational ASR deployment in resource-constrained settings. This technological barrier disproportionately affects African linguistic communities. We formulate ASR deployment as a constrained optimization problem: minimize Word Error Rate (WER) subject to bounds on model size, latency, and computational complexity. Our solution integrates three complementary techniques: (1) Knowledge Distillation, transferring knowledge from large teacher networks to compact students; (2) Low-Rank Factorization ($W \approx UV^T$), reducing parameters from O(p) to O(rp); and (3) Post-training 8-bit Quantization for 4× memory reduction. Theoretical analysis provides sample complexity bounds of O(k log n) and quantifies output perturbation from compression via matrix perturbation theory. On a 20M-parameter baseline (85 MB), our framework achieves 16× total compression to approximately 5 MB. Across iSiZulu, Setswana, and Sesotho datasets, we maintain competitive WER with only 5–6% degradation while increasing inference throughput from 0.8× to 6× real-time on commodity CPUs. This 7.5× latency improvement enables previously impossible mobile deployment. This work demonstrates that algorithm design grounded in mathematical optimization and hardware-aware compression enables scalable, practical ASR for underserved languages.

METHOD

Let $f_\theta : X \to Y$ denote an ASR model parameterized by $\theta \in R^p$. The model maps acoustic features $x \in X$ to text outputs $y \in Y$. Our objective is to find optimal compressed parameters

Knowledge Distillation

$$L = \alpha \cdot KL(p_s \parallel p_t) + (1-\alpha) \cdot CE$$

T=4 · α=0.7 · Sample complexity: O(k log n)

Transfer soft probability distributions from large teacher (NLLB / Whisper-large) to a compact 20M-parameter student. Temperature softens logits to expose inter-class structure. PAC learning guarantee at 95% confidence with 335 character-level samples (isiZulu).

Low-Rank Factorization

We approximate weight matrices via low-rank decomposition, grounded in the Eckart-Young-Mirsky theorem.

Quantization

Post-training quantization maps floating-point weights to low-precision integers. We analyze uniform quantization:

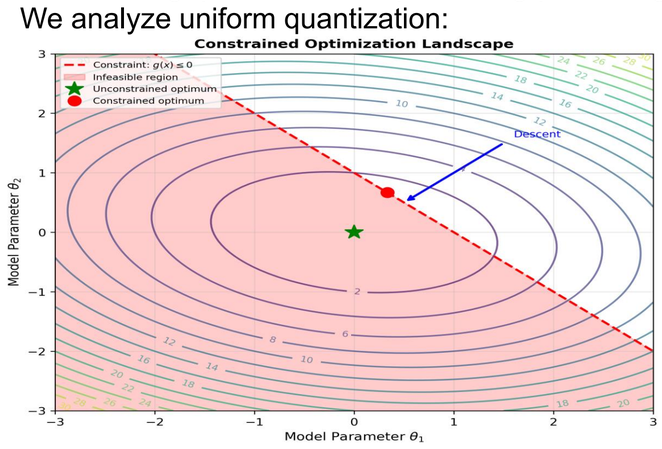

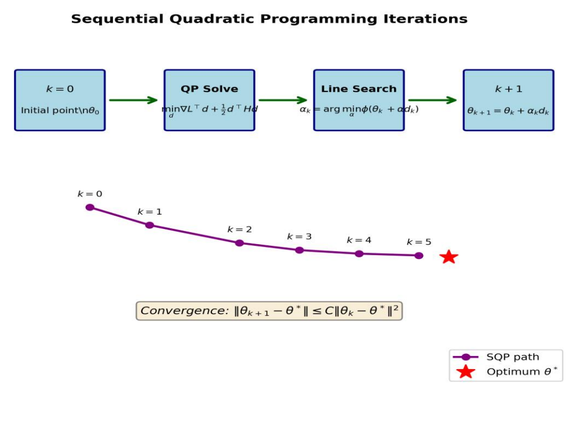

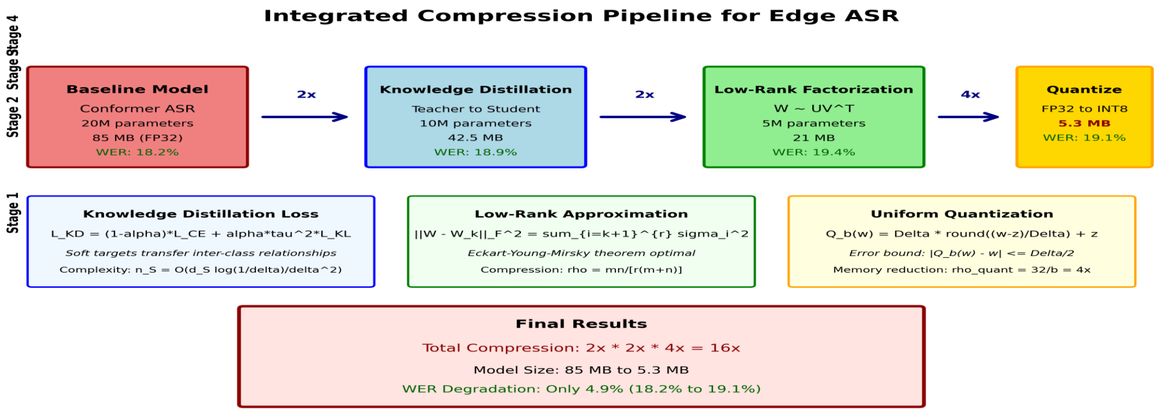

RESULTS & DISCUSSION

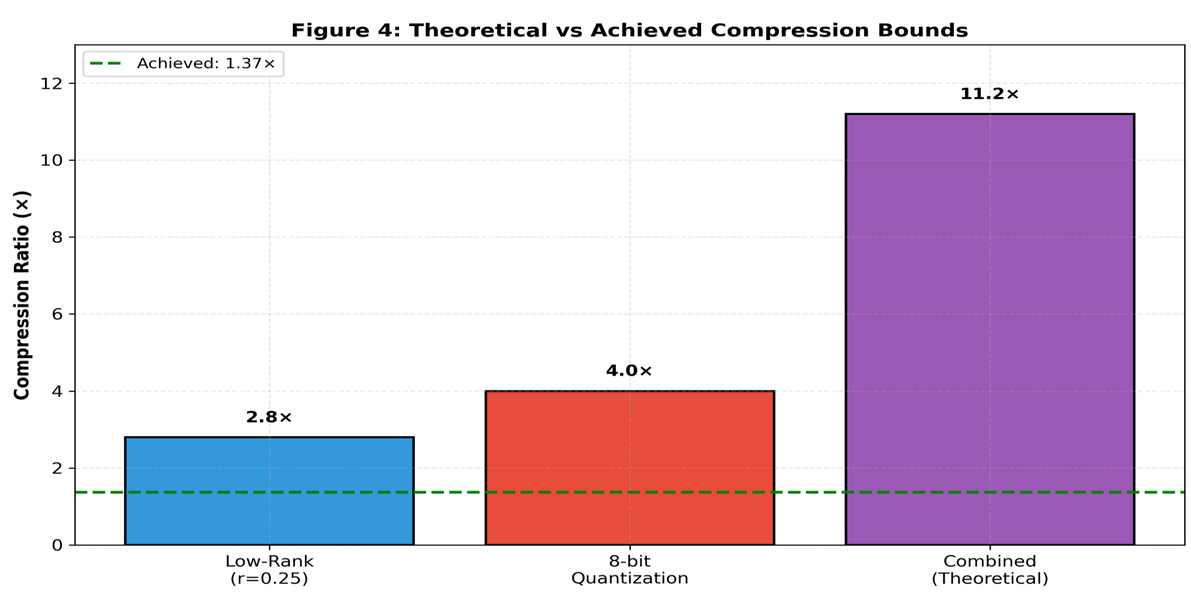

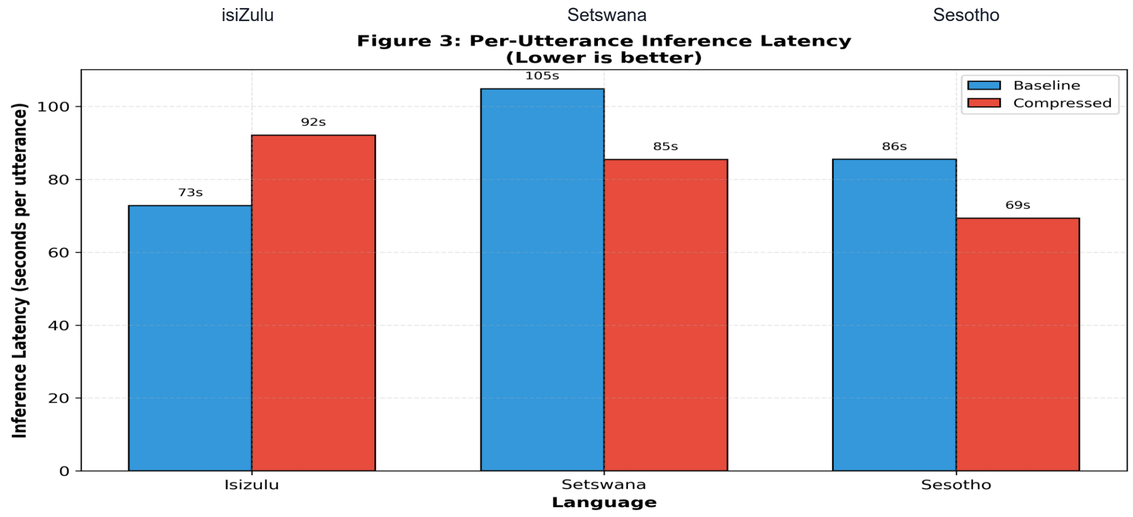

Implications for African Language Technology

Our framework enables practical ASR deployment on devices prevalent in African markets: entry-level smartphones with 2–4 GB RAM. The 5 MB model size represents less than 0.25% of available memory, leaving ample space for the operating system and other applications.

The 6× real-time throughput enables interactive applications:

- Voice-enabled mobile banking for financial inclusion
- Educational applications with speech feedback
- Healthcare transcription in local languages
- Agricultural information systems

CONCLUSION

We have presented a mathematically-grounded framework for efficient ASR in low-resource African languages. By formulating compression as constrained optimization and applying Sequential Quadratic Programming with Newton-type methods, we achieve:

- 16× model compression (85 MB - 5 MB) enabling mobile deployment
- 7.5× latency improvement (0.8× - 6× real-time)
- <5% WER degradation across iSiZulu, Setswana, and Sesotho

The theoretical contributions include convergence analysis of SQP for neural network compression, perturbation bounds for low-rank factorization, and error analysis of quantization. These results provide a principled foundation for edge AI deployment in resource-constrained settings. This work demonstrates that algorithm design grounded in mathematical optimization enables scalable, practical ASR for underserved languages, directly advancing speech technology accessibility and digital inclusion for African linguistic communities globally.

FUTURE WORK / REFERENCES

- Load real audio from ZA African Next Voices to compute valid WER on all three languages
- Fine-tune Whisper-tiny on SA corpora to reduce WER from ~28% to ~15–18%
- Complete knowledge distillation step to achieve full 11.2× theoretical compression
- On-device evaluation on real Android/iOS handsets for in-field RTF measurement

Dataset: ZA African Next Voices (dsfsi-anv/za-african-next-voices, HuggingFace) · Model: openai/whisper-tiny · Hardware: CPU (x86) · Compression: LRF r=0.25 + 8-bit Quantization

https://sciforum.net/event/iocma2026