

FlashSpeech vs. Diffusion-Based Models in Zero-Shot Speaker Verification Latency

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: How does FlashSpeech’s inference latency compare to diffusion-based baselines on zero-shot speaker verification tasks across diverse acoustic environments. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: FlashSpeech: Efficient Zero-Shot Speech Synthesis. Research question: How does FlashSpeech’s inference latency compare to diffusion-based baselines on zero-shot speaker verification tasks across diverse acoustic environments?.

2 Methodology

Systematic literature search across multiple databases yielded 7 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

3 Results

7 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.5/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

References

- <http://arxiv.org/abs/2002.03562v2>
- <http://arxiv.org/abs/2508.18913v1>
- <http://arxiv.org/abs/2404.14700v4>