

FlashSpeech Quality-Efficiency Trade-offs Against VALL-E and AudioGen on LJ Speech

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: What is the trade-off between the quality (e.g., MOS score) and inference efficiency of FlashSpeech when compared to VALL-E and AudioGen on the LJ Speech benchmark. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.7/10. 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: What is the trade-off between the quality (e.g., MOS score) and inference efficiency of FlashSpeech when compared to VALL-E and AudioGen on the LJ Speech benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.7/10.

3 Results

10 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 4.7/10.

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/2307.10550v1>
- <http://arxiv.org/abs/2411.18222v1>
- <http://arxiv.org/abs/2404.14700v4>