Qwen32Bq4 Performance on MMSU Subsets: Emotional vs Neutral Speech Fine-Tuning Effects

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does Qwen32Bq4's performance on language understanding tasks vary when fine-tuned on different subsets of the MMSU benchmark (e.g., emotional speech vs. neutral speech), measured by F1 scores. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: MMSU: A Massive Multi-task Spoken Language Understanding and Reasoning Benchmark. Research question: How does Qwen32Bq4's performance on language understanding tasks vary when fine-tuned on different subsets of the MMSU benchmark (e.g., emotional speech vs. neutral speech), measured by F1 scores?.

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 3.3/10.

3 Results

10 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.3/10.

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/2506.04779v3
- http://arxiv.org/abs/2502.12672v4
- http://arxiv.org/abs/2207.08179v1