MobileVLM Throughput Latency Scaling on Heterogeneous Mobile Hardware

Assignee Research

June 5, 2026

Abstract

This report synthesises findings from 5 peer-reviewed papers addressing the following research question: How does the throughput latency of MobileVLM's efficient projector architecture scale when deployed on heterogeneous mobile hardware compared to standard transformer-based projectors. 9 claims were extracted from source literature; 9 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 9.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: MobileVLM : A Fast, Strong and Open Vision Language Assistant for Mobile Devices. Research question: How does the throughput latency of MobileVLM's efficient projector architecture scale when deployed on heterogeneous mobile hardware compared to standard transformer-based projectors?.

2 Methodology

Systematic literature search across multiple databases yielded 5 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 9.2/10.

3 Results

5 papers retrieved. 9 claims extracted; 9 independently verified. Quality review score: 9.2/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
MobileVLM is a multimodal vision language model (MMVLM) designed to run on mobile devices.	✓	0.29
MobileVLM combines various architectural designs and techniques optimized for mobile devices.	✓	0.16
MobileVLM includes language models with 1.4B and 2.7B parameters, trained from scratch.	✓	0.22
MobileVLM uses a multimodal vision model pretrained in the CLIP fashion.	✓	0.26
MobileVLM achieves cross-modality interaction via an efficient projector.	✓	0.19
MobileVLM demonstrates performance on par with much larger models on several typical VLM benchmarks.	✓	0.18
MobileVLM achieves state-of-the-art inference speed of 21.5 tokens per second on a Qualcomm Snapdragon 888 CPU.	✓	0.25
MobileVLM achieves state-of-the-art inference speed of 65.3 tokens per second on an NVIDIA Jeston Orin GPU.	✓	0.25
The code for MobileVLM will be made available at https://github.com/Meituan-AutoML/MobileVLM.	✓	0.25

References

- https://doi.org/10.48550/arxiv.2312.16886
- https://doi.org/10.48550/arxiv.2404.16112
- https://doi.org/10.48550/arxiv.2409.02889