

Multimodal Model Scaling with Structural Graph Priors vs. Pure Attention on Flickr30k Entities

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does the scaling behavior of multimodal models differ when using structural graph priors versus pure attention, as measured by BLEU score improvements on the Flickr30k Entities benchmark under. Background: AI-driven prediction algorithms have the potential to enhance emergency medicine by enabling rapid and accurate decision-making regarding patient status and potential deterioration. However, the integration of multimodal data, including raw waveform signals, remains. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Enhancing clinical decision support with physiological waveforms – a multimodal benchmark in emergency care. Research question: How does the scaling behavior of multimodal models differ when using structural graph priors versus pure attention, as measured by BLEU score improvements on the Flickr30k Entities benchmark under low-noise conditions?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

13 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 4.5/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
The multimodal model integrating ECG waveforms and clinical routine data achieved a macro-AUROC of 0.8256 (0.8222, 0.828	×	0.08
The multimodal model integrating ECG waveforms and clinical routine data achieved an AUROC of 0.9115 (0.8991, 0.9222) fo	×	0.08
The model’s predictive performance ranges from an AUROC of 0.7405 for the musculoskeletal system and connective tissue (×	0.03
The model predicted 609 out of 1,428 individual ICD diagnoses with high accuracy, defined as conditions where the lower	×	0.05
The inclusion of ECG waveforms alongside clinical routine data consistently improves performance over ECG features and c	×	0.05
The highest gains are observed for XII (Skin, 13.06%), VI (Nervous, 10.75%), and XIX (Injury, Poisoning, 9.95%).	×	0.01
Smaller improvements are seen for III (Blood, 2.36%) and IV (Endocrine, 5.25%).	×	0.01
For clinical deterioration, the model achieves an AUROC of 0.9070.	×	0.11
For ICU admissions, the model reports an overall AUROC of 0.9063.	×	0.03
For mortality predictions, the model exhibits an overall AUROC of 0.9168.	×	0.03
The MDS-ED pipeline involves feature collection encompassing patient demographics, biometrics, vital parameters and tren	×	0.11
The MDS-ED pipeline predicts patient discharge diagnoses out of 1428 cardiac and non-cardiac ICD10-CM codes and predicts	×	0.12
The MDS-ED pipeline collects features from a window of 90 minutes from the patient’s arrival at the ED.	×	0.01
The MDS-ED was created by linking ECG waveforms from the MIMIC-IV-ECG dataset to clinical features and outcomes from the	×	0.03

References

- <http://arxiv.org/abs/2403.09832v1>
- <http://arxiv.org/abs/2407.17856v4>
- <http://arxiv.org/abs/2507.21100v1>