Contrastive Loss Functions and Adversarial Robustness in ECG Foundation Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 9 peer-reviewed papers addressing the following research question: How does the choice of contrastive loss function in self-supervised pretraining affect the robustness of ECG foundation models under adversarial perturbations, measured by classification accuracy on. 10 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: CLEF: Clinically-Guided Contrastive Learning for Electrocardiogram Foundation Models. Research question: How does the choice of contrastive loss function in self-supervised pretraining affect the robustness of ECG foundation models under adversarial perturbations, measured by classification accuracy on MIT-BIH and PTB-XL datasets?.

2 Methodology

Systematic literature search across multiple databases yielded 9 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.0/10.

3 Results

9 papers retrieved. 10 claims extracted; 0 independently verified. Quality review score: 4.0/10.

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
CLEF-M achieves an AUROC of 0.8194 on downstream classification tasks.	×	0.10
MOCO (2020) achieves an AUROC of 0.8112 on downstream classification tasks.	×	0.04
SimCLR (2020) achieves an AUROC of 0.8243 on downstream classification tasks.	×	0.04
BYOL (2020) achieves an AUROC of 0.8221 on downstream classification tasks.	×	0.04
The pretraining dataset includes 10-second 12-lead ECG samples.	×	0.11
The negative weighting loss (Lw) is calculated as (1+fsim(zi, zj))/2 * (1-Wi,j).	×	0.01
The dissimilarity alignment loss (Ld) is calculated as (1+fsim(zi, zk))/2 * (1-Wi,k).	×	0.00
Clinical metadata includes age, sex, and blood pressure.	×	0.06
The model uses clinically validated risk scores for pretraining.	×	0.06
The model is pretrained on 12-lead ECGs to learn generalizable representations.	×	0.12

References

- http://arxiv.org/abs/2512.02180v1
- http://arxiv.org/abs/2004.13701v1
- http://arxiv.org/abs/2103.12676v2