

# Pretraining Objectives and Adversarial Robustness in ECG Foundation Models

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: Do different pretraining objectives influence the robustness of ECG foundation models against adversarial noise when measured by classification accuracy degradation. 10 claims were extracted from source literature; 4 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Self-supervised representation learning from 12-lead ECG data. Research question: Do different pretraining objectives influence the robustness of ECG foundation models against adversarial noise when measured by classification accuracy degradation?.

## 2 Methodology

Systematic literature search across multiple databases yielded 16 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.0/10.

## 3 Results

16 papers retrieved. 10 claims extracted; 4 independently verified. Quality review score: 6.0/10.

## 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
The best-performing methods in speech representation learning rely on latent forecasting tasks.	×	0.12
SimCLR uses a noise contrastive loss to attract positive copies and repel negative instances in the batch.	×	0.09
BYOL does not explicitly rely on contrasting against negative samples in the same batch.	×	0.06
SwAV relies on contrasting cluster assignments rather than individual instances.	×	0.00
The xresnet1d50 model architecture is used, which performs comparably to the best-performing xresnet1d101 within error b	×	0.02
The linear evaluation performance of the best-performing method (adaptation of contrastive predictive coding) is only 0.	✓	0.31
Fine-tuned models show improvements in downstream performance of roughly 1% compared to supervised performance.	✓	0.22
The work establishes the feasibility of extracting discriminative representations from ECG data via self-supervised learning	✓	0.31
The work shows improvements in label efficiency and robustness against physiological noise.	✓	0.15
The work is the first comprehensive assessment of its kind in the ECG domain carried out exclusively on publicly available data	×	0.11

## References

- <http://arxiv.org/abs/2512.02180v1>

- <http://arxiv.org/abs/2605.12241v1>
- <http://arxiv.org/abs/2103.12676v2>