

Synthetic Pretraining Performance in Tabular Foundation Models vs. Real-Data Benchmarks

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: Can synthetic pretraining strategies for tabular foundation models match the performance of large-scale real-data pretraining on standard structured data benchmarks. 13 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.1/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Pretraining Strategies and Scaling for ECG Foundation Models: A Systematic Study. Research question: Can synthetic pretraining strategies for tabular foundation models match the performance of large-scale real-data pretraining on standard structured data benchmarks?.

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 4.1/10.

3 Results

16 papers retrieved. 13 claims extracted; 0 independently verified. Quality review score: 4.1/10.

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 study covers five different pretraining methodologies trained on a corpus comprising over 11 million ECG samples.	×	0.09
State space models (SSM) are confirmed as the superior architecture choice across all pretraining paradigms compared to	×	0.09
The CPC pretraining strategy shows the strongest and most transferable representations across diverse clinical tasks.	×	0.15
The data2vec pretraining strategy consistently lags behind other methodologies across all evaluation modes and scaling r	×	0.08
Scaling behavior is most clearly identified for the CPC and JEPA pretraining methodologies.	×	0.14
Lower pretraining loss correlates with small residual errors in downstream tasks.	×	0.05
The S4 backbone with a model dimension of 512 consistently outperforms larger dimensions (768, 1024) and alternative bac	×	0.01
All models in the study operate at a sampling rate of 240 Hz on 12-lead ECG inputs.	×	0.03
The default backbone adopted for the study is a 4-layer S4 model with a dimension of 512.	×	0.02
The CNN stem used in all models consists of four convolutional layers with batch normalization.	×	0.03
The Transformer backbone variant utilizes RoPE positional encoding and GELU activations.	×	0.03
Recent ECG foundation models have been pre-trained on tens of millions of recordings.	×	0.11
Structured state space models have previously shown superior performance on long sequences in supervised ECG settings co	×	0.14

References

- <http://arxiv.org/abs/2412.10104v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2605.12241v1>