

Alignment Techniques and Inference Efficiency in Contrastive Time-Series Foundation Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What is the impact of alignment techniques on the inference efficiency of contrastive time-series foundation models when deployed in resource-constrained environments, as measured by. 17 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: What Constitutes Good Contrastive Learning in Time-Series Forecasting?. Research question: What is the impact of alignment techniques on the inference efficiency of contrastive time-series foundation models when deployed in resource-constrained environments, as measured by latency-throughput trade-offs?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

15 papers retrieved. 17 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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 ETT datasets consist of data from six power load features and measurements of oil temperature.	×	0.04
The ECL dataset contains the electricity consumption of 312 clients.	×	0.04
The ECL dataset was converted into hourly-level measurements, following previous work (Yue et al., 2022).	×	0.04
The LSTM-based encoder includes five unidirectional LSTM layers with a hidden dimension of 320 in each layer.	×	0.02
The LSTM-based encoder architecture has a total of 660K learnable parameters.	×	0.01
The TCN-based encoder consists of 10 dilated convolution layers with a hidden dimension of 320 and a kernel size of 3	×	0.04
The TCN-based encoder architecture has a total of 637K learnable parameters.	×	0.01
The Transformer-based encoder contains 5 In-former layers with a hidden size of 128 and 8 attention heads.	×	0.03
The Transformer-based encoder architecture has a total of 655K learnable parameters.	×	0.01
All three models utilize a linear layer as the first embedding layer to map the input features into hidden dimensions.	×	0.04
All three models are trained with a peak learning rate of 0.001.	×	0.04
A cosine learning rate scheduler is used for MoCo2-framework models.	×	0.04
Each model trains for 30 epochs with early stopping based on the Dev performance to prevent overfitting for end-to-end t	×	0.06
For two-step learning, an encoder is trained on the training set with SSCL for 600 training iterations.	×	0.07
When fine-tuning a pre-trained encoder, the same training hyper-parameters as in end-to-end training are used.	×	0.07
The MSE is used as the main evaluation metric, while MAE results are also reported.	×	0.04
The overall performance of each model is represented by averaging the MSE across all prediction lengths and the three da	×	0.03

References

- <http://arxiv.org/abs/2306.12086v2>
- <http://arxiv.org/abs/2303.00106v1>
- <http://arxiv.org/abs/2109.05472v2>