Pretraining Computational Cost and Cross-Domain Generalization in Tabular Foundation Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: What is the relationship between pretraining computational cost (training time) and cross-domain generalization of tabular foundation models, evaluated using a standardized set of benchmarks such as. 11 claims were extracted from source literature; 7 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 7.4/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Amazon SageMaker Autopilot. Research question: What is the relationship between pretraining computational cost (training time) and cross-domain generalization of tabular foundation models, evaluated using a standardized set of benchmarks such as OpenML-2019 and Amazon Employee Access datasets?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.4/10.

3 Results

4 papers retrieved. 11 claims extracted; 7 independently verified. Quality review score: 7.4/10.

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
Amazon SageMaker Autopilot is a fully managed system that provides an automatic machine learning solution.	✓	0.36
Given a tabular dataset and the target column name, Autopilot identifies the problem type.	✓	0.25
Autopilot analyzes the input data and produces a diverse set of complete ML pipelines.	✓	0.22
Autopilot tunes generated pipelines to produce a leaderboard of candidate models for user selection.	×	0.10
The diversity of generated pipelines allows users to balance trade-offs such as model accuracy versus latency.	×	0.15
Autopilot exposes the training pipelines in addition to the final models, allowing users to customize the generated trai	✓	0.17
Autopilot is designed to serve users with different levels of expertise by being neither a fully black-box solution nor	✓	0.22
The infrastructure choices in Autopilot allow for scalability.	×	0.14
The infrastructure choices in Autopilot enable the generation of high quality models.	×	0.12
Autopilot supports the consumption of artifacts from offline meta-learning.	✓	0.17
Autopilot integrates with the SageMaker system to allow trained models to be used in a production setting.	✓	0.20

References

- https://doi.org/10.48550/arxiv.2305.02997

- https://doi.org/10.18653/v1/2023.emnlp-main.917
- https://doi.org/10.1145/3399579.3399870