

Scalability of Contrastive Self-Supervised Learning for Tabular Data via Throughput and Representation Quality

Assignee Research

June 11, 2026

Abstract

Abstract Tabular data, spreadsheets organized in rows and columns, are ubiquitous across scientific fields, from biomedicine to particle physics to economics and climate science ^{1,2}. The fundamental prediction task of filling in missing values of a label column based on the rest of the columns is essential for various applications as diverse as biomedical risk models, drug discovery and materials science. Although deep learning has revolutionized learning from raw data and led to numerous high-profile success stories ^{3–5}, gradient-boosted decision trees ^{6–9} have dominated tabular data for th

1 Introduction

This paper examines: Accurate predictions on small data with a tabular foundation model. Research question: How scalable are contrastive self-supervised learning methods for tabular data when applied to larger datasets, as measured by training throughput and representation quality on benchmarks like TabNet or XGBoost?.

2 Methodology

Systematic literature search across multiple databases yielded 12 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

12 papers retrieved. 8 claims extracted; 6 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
Tabular data are ubiquitous across scientific fields, from biomedicine to particle physics to economics and climate science	✓	0.28
Gradient-boosted decision trees have dominated tabular data for the past 20 years.	✓	0.26
TabPFN outperforms all previous methods on datasets with up to 10,000 samples by a wide margin.	✓	0.25
TabPFN uses substantially less training time than previous methods.	×	0.12
In a classification setting, TabPFN achieves superior performance in 2.8 seconds compared to an ensemble of the strongest	×	0.12
TabPFN is a generative transformer-based foundation model.	✓	0.22
TabPFN allows fine-tuning, data generation, density estimation, and learning reusable embeddings.	✓	0.24
TabPFN is a learning algorithm that is itself learned across millions of synthetic datasets.	✓	0.22

References

- <https://doi.org/10.48550/arxiv.2310.18541>
- <https://doi.org/10.26434/chemrxiv-2022-jjm0j-v4>
- <https://doi.org/10.1038/s41586-024-08328-6>