Pretraining Sample Diversity and Robustness to Column Permutation Noise in Tabular Few-Shot Learning

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: What is the correlation between pretraining sample diversity and robustness to column permutation noise in tabular few-shot learning evaluations. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Tabular Few-Shot Generalization Across Heterogeneous Feature Spaces. Research question: What is the correlation between pretraining sample diversity and robustness to column permutation noise in tabular few-shot learning evaluations?.

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 4.2/10.

3 Results

12 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 4.2/10.

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 uses 118 tabular classification datasets from the UCI Machine Learning Repository.	×	0.07
65 of the 118 datasets have more than two prediction classes and were binarized for the study.	×	0.03
FLAT models are trained and tested using an N-fold evaluation procedure.	×	0.04
Each fold is used once as the testing collection Dtest, while the remaining N -1 folds form Dtrain.	×	0.02
Feature columns are standardized to mean 0 and variance 1.	×	0.05
During training, feature columns are randomly subsampled for both Dmeta and Dtarget as a form of data augmentation.	×	0.03
FLAT results are averaged over multiple random seeds.	×	0.03
The study employs a randomized sampling procedure for imbalanced few-shot learning, with the number of positive examples	×	0.07
Dmeta contains at least one example of each class unless N meta = 1.	×	0.04
The study compares standard K-shot and binomial sampling approaches in Appendix A.4.1.	×	0.04
The study evaluates FLAT against several baselines including logistic regression (LR), k-nearest neighbors (KNN), suppor	×	0.02
Iwata is meta-trained and tested using the same N-fold evaluation procedure as FLAT.	×	0.02
All remaining baselines require a training dataset with the same feature space as the test dataset.	×	0.03
The study does not compare against TabLLM since the setup does not assume access to semantically meaningful columns.	×	0.01

References

- `http://arxiv.org/abs/2311.10051v1`

- http://arxiv.org/abs/2603.10254v1
- http://arxiv.org/abs/2311.14544v1