Adversarial Tabular Foundation Models Outperform Traditional ML on Cross-Domain Benchmarks

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: Do tabular foundation models trained with adversarial objectives maintain higher accuracy than traditional ML baselines on cross-domain tabular 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 3.1/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: ExPLoRA: Parameter-Efficient Extended Pre-Training to Adapt Vision Transformers under Domain Shifts. Research question: Do tabular foundation models trained with adversarial objectives maintain higher accuracy than traditional ML baselines on cross-domain tabular 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 3.1/10.

3 Results

16 papers retrieved. 13 claims extracted; 0 independently verified. Quality review score: 3.1/10.

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
ExPLoRA achieves a top-1 accuracy of 79.3% on the fMoW-RGB benchmark.	×	0.03
ExPLoRA improves top-1 accuracy by 1.5% over previous methods on the fMoW-RGB benchmark.	×	0.03
ExPLoRA uses 6% of the ViT encoder parameters compared to full pre-training.	×	0.12
ExPLoRA requires 100 GPU hours for training on the fMoW-RGB benchmark.	×	0.02
Full pre-training on the fMoW-RGB benchmark requires over 960 GPU hours.	×	0.08
ExPLoRA achieves an 8.2% improvement in linear probing accuracy on the fMoW dataset compared to baselines.	×	0.10
ExPLoRA requires 8x to 10x less compute than fully-pretrained prior state-of-the-art methods on satellite datasets.	×	0.07
ExPLoRA uses 16x fewer trainable parameters than fully-pretrained prior state-of-the-art methods on satellite datasets.	×	0.08
The fMoW dataset contains high-resolution satellite images paired with one of 62 classification labels.	×	0.05
The D-ExPLoRA-[L]-r64 configuration achieves an average accuracy of 79.28% on the fMoW-RGB benchmark.	×	0.02
ExPLoRA-initializations with LoRA fine-tuning outperform other unsupervised initializations paired with PEFT techniques	×	0.14
The ExPLoRA-[L]-r64-itr200k configuration utilizes 18.7M total parameters and 0.8M trainable parameters.	×	0.03
The DinoV2-[0, 1, L, L 1]-r64 configuration utilizes 55.8M trainable parameters.	×	0.04

References

- http://arxiv.org/abs/2512.03307v1
- http://arxiv.org/abs/2406.10973v5

- http://arxiv.org/abs/2402.01204v4