

# Adversarial Training and Label Noise Robustness in Tabular Foundation Models

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: To what extent do adversarially trained tabular foundation models maintain robustness against label noise (e.g., 20% random label corruption) compared to standard models on benchmarks like TabMNAR,. 10 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Mitigating Noisy Supervision Using Synthetic Samples with Soft Labels. Research question: To what extent do adversarially trained tabular foundation models maintain robustness against label noise (e.g., 20% random label corruption) compared to standard models on benchmarks like TabMNAR, measured by F1-score degradation?.

## 2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.7/10.

## 3 Results

14 papers retrieved. 10 claims extracted; 1 independently verified. Quality review score: 4.7/10.

## 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 authors Y. Lu and W. He are affiliated with the Department of Computing and Software at McMaster University, Hamilto	×	0.00
The manuscript was received on May 6, 2022.	×	0.00
The paper extends the work cited as reference [8] to combat the negative influence of noisy labels.	×	0.08
The proposed method generates synthetic samples by searching for K nearest neighbours based on learned representations a	✓	0.18
In the example provided in Figure 1, averaging three images (two dogs, one cat) results in a synthetic label with a prob	×	0.03
The proposed method assigns dynamic weights to selected samples and convexly combines them to generate synthetic samples	×	0.08
When trained with noisy labels, overparameterized DNNs first fit training data with clean labels during an early learnin	×	0.14
Recent study [27] has theoretically proven that the early learning phenomenon occurs in a simple linear model.	×	0.05
The gradient of cross entropy loss with respect to parameters $\Theta$ is defined as the negative average sum of the Jacobian m	×	0.04
In a clean training data scenario, the term $(\pi - y_i)$ for the true class entry is always negative, while the rest of the	×	0.05

## References

- <http://arxiv.org/abs/2110.06513v2>

- <http://arxiv.org/abs/2406.16966v1>
- <http://arxiv.org/abs/2512.03307v1>