

Adversarial Training Effects on Graph Neural Network Inference Throughput in Spatio-Temporal Data Streams

Assignee Research

June 1, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does adversarial training impact the inference throughput of graph neural networks on high-frequency spatio-temporal data streams compared to standard training. Missing data imputation (MDI) is a fundamental problem in many scientific disciplines. Popular methods for MDI use global statistics computed from the entire data set (e.g., the feature-wise medians), or build predictive models operating independently on every instance. 17 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Missing Data Imputation with Adversarially-trained Graph Convolutional Networks. Research question: How does adversarial training impact the inference throughput of graph neural networks on high-frequency spatio-temporal data streams compared to standard training?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

15 papers retrieved. 17 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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 GINN framework was evaluated on 20 real-world datasets from the UCI Machine Learning Repository.	×	0.13
The datasets used for evaluation contain categorical, numerical, and mixed datasets, ranging from 150 observations to 30	×	0.04
Every dataset is divided into training (70%) and test (30%) sets.	×	0.03
Missingness is introduced completely at random on the training set with 4 different levels of noise: 10%, 20%, 30%, and	×	0.02
The evaluation focuses on imputation performance and the accuracy of post-imputation prediction.	×	0.07
An embedding dimension of the hidden layer of 128 is used, sufficient for an overcomplete representation for all the dat	×	0.02
The model is trained for a maximum of 10000 iterations with an early stopping strategy for the reconstruction loss over	×	0.03
The critic used is a simple 3-layer feed-forward network trained 5 times for each optimization step of the autoencoder.	×	0.03
The Adam optimizer is used for both networks with a learning rate of 1×10^{-3} and 1×10^{-5} respectively for autoencoder an	×	0.04
All experiments are repeated five times and average performance and standard deviation are collected.	×	0.02
The code for replicating the experiments and using the GINN algorithm is released as an open-source library on the web.	×	0.04
The characteristics of the 20 datasets are summarized in Table 1.	×	0.04
The performance of the algorithm is evaluated on three real-world datasets with pre-existing missing values.	×	0.10
The computational cost of the method is compared to other state-of-the-art approaches.	×	0.07
In one case, missing values are also present in the training labels, and the method is applied to a semi-supervised scen	×	0.05
The median and standard deviation values for various datasets and missing percentages are provided in Table (p19).	×	0.03
The count of wins for different methods across various missing percentages is provided in Table (p20).	×	0.04

References

- <http://arxiv.org/abs/2502.04140v2>
- <http://arxiv.org/abs/1905.01907v2>
- <http://arxiv.org/abs/2406.00552v4>