Fine-Tuning Llama-3.1-8B on Synthetic Obfuscation Patterns for Generalization to Unseen Vulnerabilities in Big-Vul

Assignee Research

June 13, 2026

Abstract

Amyotrophic Lateral Sclerosis (ALS) is a rare neurodegenerative disease, and high-quality EEG data from ALS patients are scarce. This data scarcity, coupled with severe class imbalance between ALS and healthy control recordings, poses a challenge for training reliable machine learning classifiers. In this work, we address these issues by generating synthetic EEG signals for ALS patients using a Conditional Wasserstein Generative Adversarial Network (CWGAN). We train CWGAN on a private EEG dataset (ALS vs. non-ALS) to learn the distribution of ALS EEG signals and produce realistic synthetic sam

1 Introduction

This paper examines: Synthetic ALS-EEG Data Augmentation for ALS Diagnosis Using Conditional WGAN with Weight Clipping. Research question: To what extent does fine-tuning Llama-3.1-8B on synthetic obfuscation patterns improve generalization to unseen vulnerability types in the Big-Vul dataset compared to standard fine-tuning?.

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 7.4/10.

3 Results

16 papers retrieved. 9 claims extracted; 7 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
The synthetic ALS EEG signals generated by the WGAN exhibit oscillatory patterns and amplitude dynamics that fall within	✓	0.24
The synthetic ALS signals do not appear as simple repeats of training data; they show varied patterns, indicating that t	✓	0.28
Clinicians and EEG experts found the synthetic ALS EEG segments to be indistinguishable from real ALS EEG segments in ma	✓	0.23
The critic's loss settles into a narrow oscillation around a stable mean, while the generator's loss steadily declines t	✓	0.28
Weight clipping to [-0.01, +0.01] effectively enforced the 1-Lipschitz constraint and promoted stable updates.	✓	0.35
The EEGET-ALS Dataset contains raw 32-channel EEG recordings sampled at 256 Hz from six ALS patients and 170 healthy con	✓	0.26
The EEGET-ALS Dataset has a representation of $\approx$3.4% ALS samples.	×	0.13
Each EEG segment is linearly rescaled into the [–1, 1] range using the formula (x - x.min()) / (x.max() - x.min()) * 2 –	×	0.11
Uniform scaling is applied file-wise to ensure that the GAN sees inputs on a consistent scale, which is critical for sta	✓	0.29

References

- http://arxiv.org/abs/2512.09006v1
- http://arxiv.org/abs/2601.08691v1

- http://arxiv.org/abs/2506.16243v1