Synthetic Feature Dimensionality and Robustness of Self-Supervised Representations Under Noise

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: To what extent does increasing synthetic tabular feature dimensions affect the robustness of self-supervised representations against noise injection in downstream classification tasks. 17 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Elastic Weight Consolidation Improves the Robustness of Self-Supervised Learning Methods under Transfer. Research question: To what extent does increasing synthetic tabular feature dimensions affect the robustness of self-supervised representations against noise injection in downstream classification tasks?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

11 papers retrieved. 17 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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 FIM is computed using the first 10000 images of ImageNet1k for a pre-trained ImageNet1k DINO model.	×	0.05
Three downstream tasks are considered: Waterbirds, Celeb-A, and CIFAR10.	×	0.10
Celeb-A and Waterbirds are datasets with known or constructed biases used for analyzing worst sub-group performance.	×	0.09
The deterioration of worst sub-group performance is interpreted as a catastrophic forgetting event for a robust SSL mode	×	0.09
EWC is leveraged to minimize the performance penalty due to catastrophic forgetting.	×	0.02
The FIM analysis involves fine-tuning a pre-trained DINO ViT-B/16 ImageNet1k model on the CIFAR10 dataset.	✓	0.17
Performance is evaluated by computing top-1 accuracy on the CIFAR10 dataset and top-1 accuracy on the ImageNet1k dataset	×	0.03
Reverse transfer performance from CIFAR10 to ImageNet1k is evaluated by attaching the ImageNet1k classification head fro	×	0.10
As the regularization weight is increased, both methods improve in terms of their performance on ImageNet1k.	×	0.06
FIM regularization does not decrease performance on CIFAR10, unlike the naive L2 regularization approach.	×	0.03
The L2 regularization approach manages to fully recover 77% top-1 accuracy on ImageNet1k.	×	0.02
FIM regularization saturates at around 73% top-1 accuracy on ImageNet1k, even for very high regularization weights.	×	0.03
To keep fine-tuned representations close to their initial SSL values, techniques used in CL, particularly EWC, are consi	×	0.10
EWC regularizes model parameters towards their optimal values on previous tasks using the Fisher Information Matrix (FIM	×	0.11
The regularization is based on the Laplace approximation.	×	0.04
The posterior distribution with respect to the SSL task is approximated with a Normal distribution using a Taylor's expa	×	0.08
A point estimate to the posterior distribution is derived via Bayes rule.	×	0.03

References

- http://arxiv.org/abs/2101.09825v1
- http://arxiv.org/abs/1906.11951v1
- http://arxiv.org/abs/2210.16365v1