Diffusion-Driven Input Adaptation for Cross-Domain Streaming Data Performance

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: Can diffusion-driven input adaptation maintain consistent performance gains over traditional test-time training when evaluated on streaming data with varying entropy levels in cross-domain scenarios. 15 claims were extracted from source literature; 2 were 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: Everything to the Synthetic: Diffusion-driven Test-time Adaptation via Synthetic-Domain Alignment. Research question: Can diffusion-driven input adaptation maintain consistent performance gains over traditional test-time training when evaluated on streaming data with varying entropy levels in cross-domain scenarios?.

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 4.5/10.

3 Results

15 papers retrieved. 15 claims extracted; 2 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
Fine-tuning the source model on synthetic data generated by the Mix of Diffusion process yields a 5.5% performance impro	✓	0.17
Fine-tuning the source model on synthetic data generated by the Mix of Diffusion process yields a 6.0% performance impro	✓	0.17
Diffusion synthetic data and source data exhibit no noticeable visual differences across different timesteps t.	×	0.09
SDA consistently outperforms all baseline methods across different model architectures and sizes on ImageNet-C.	×	0.05
Compared to DDA, SDA improves accuracy by 2.5% to 2.9%.	×	0.02
Compared to GDA, SDA achieves a 2.2% improvement with ConvNeXt-T.	×	0.01
Three diffusion-driven methods (SDA, DDA, and GDA) demonstrate superior performance compared to the model adaptation met	×	0.10
DiffPure presents worse results than SDA, DDA, and GDA because it is primarily designed for adversarial attacks.	×	0.04
SDA surpasses DiffPure in all evaluated corruption types on ImageNet-C.	×	0.01
For Swin-B at timestep 500, the Source-Synthetic (Misaligned) accuracy is 61.6%, while the Synthetic-Synthetic (Aligned)	×	0.04
For ConvNeXt-B at timestep 1000, the Source-Synthetic (Misaligned) accuracy is 41.5%, while the Synthetic-Synthetic (Ali	×	0.04
On ImageNet-C, SDA achieves an average accuracy of 51.9% with ConvNeXt-B, compared to 49.4% for DDA.	×	0.02
On ImageNet-C, SDA achieves an average accuracy of 32.5% with ResNet-50.	×	0.01
Under Gaussian corruption, SDA achieves an accuracy of 60.3%, while the Source model achieves 48.0%.	×	0.04
Under JPEG corruption, SDA achieves an accuracy of 63.0%, while DiffPure achieves 34.3%.	×	0.01

References

- http://arxiv.org/abs/2409.08687v4
- http://arxiv.org/abs/2502.14293v2
- http://arxiv.org/abs/2406.04295v2