

Structural Graph Priors Enhance Zero-Shot Extraction in Noisy Multimodal Transformers

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: How does integrating structural graph priors into multimodal transformers affect zero-shot extraction accuracy on noisy image-text benchmarks like NoisyVisDial compared to pure attention baselines. Deep neural networks (DNNs) with noisy weights, which we refer to as noisy neural networks (NoisyNNs), arise from the training and inference of DNNs in the presence of noise. NoisyNNs emerge in many new applications, including the wireless transmission of DNNs, the efficient. 7 claims were extracted from source literature; 4 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Denoising Noisy Neural Networks: A Bayesian Approach with Compensation. Research question: How does integrating structural graph priors into multimodal transformers affect zero-shot extraction accuracy on noisy image-text benchmarks like NoisyVisDial compared to pure attention baselines?.

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 5.5/10.

3 Results

16 papers retrieved. 7 claims extracted; 4 independently verified. Quality review score: 5.5/10.

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
On the CIFAR-10 dataset, the denoiser outperforms the ML estimator by 4.1 dB to achieve a 60% test accuracy on the noisy	✓	0.21
On the CIFAR-10 dataset, the denoiser outperforms the ML estimator by 3 dB to achieve a 60% test accuracy on the noisy R	✓	0.20
On the CIFAR-10 dataset, the denoiser outperforms the ML estimator by 3.8 dB to achieve a 60% test accuracy on the noisy	✓	0.16
On the SST-2 dataset with the BERT model, the denoiser is 1.1 dB better than the ML estimator to achieve a 75% test accu	×	0.09
For the ShuffleNet V2 model in noisy training, the MMSEpb denoiser achieves up to 1.7 dB gains over the ML estimator to	×	0.11
For the ResNet18 model in noisy training, the MMSEpb denoiser is 13.4 dB better than the ML estimator to achieve a 60% t	✓	0.18
For the ResNet18 model in noisy training, the MMSEpb denoiser is 8.3 dB better than the ML estimator to achieve an 80% t	×	0.14

References

- <http://arxiv.org/abs/2112.10006v6>
- <http://arxiv.org/abs/2105.10699v3>

- <http://arxiv.org/abs/2505.04921v2>