Scaling Computational Throughput of Node-Based Bayesian Neural Networks vs. Traditional Ensembles

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: To what extent does the computational throughput of node-based Bayesian neural networks scale compared to traditional ensembles when applied to large-scale tabular datasets with covariate shift. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Tackling covariate shift with node-based Bayesian neural networks. Research question: To what extent does the computational throughput of node-based Bayesian neural networks scale compared to traditional ensembles when applied to large-scale tabular datasets with covariate shift?.

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 3.8/10.

3 Results

15 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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 study uses CIFAR and TINYIMAGENET datasets with corrupted test set versions provided by Hendrycks & Dietterich (2019	×	0.02
The architectures used in the experiments are VGG16, RESNET18, and PREACTRESNET18.	×	0.03
Three latent variable structures were tested: 'in' (input only), 'out' (output only), and 'both' (input and output).	×	0.03
The number of Gaussian components (K) in the variational posterior was set to 1, 2, or 4.	×	0.09
On OOD data, the optimal performance of using both input and output latent variables is similar to using only output lat	×	0.08
On OOD data, using only input latent variables produces slightly worse optimal performance compared to using output or b	×	0.08
The optimal gamma (γ) is lower when the model uses both types of latent variables (z, s) compared to other configuration	×	0.05
In Figure 8 experiments using RESNET18 on CIFAR10 with K=4 components, the NLL of noisy labels increases much faster tha	×	0.02
Higher γ prevents the model from memorizing random labels even when the majority of training labels are wrong (80% noise	×	0.06
High γ leads to improved performance on clean test sets because it prevents learning from noisy labels.	×	0.06
Figure 3 shows that variational entropy H[q(Z)] decreases over time during training.	×	0.04
The model with higher entropy (M32) performs better than the model with lower entropy (M16) across all corruption levels	×	0.03
The parameter λ controls the severity of the generated implicit corruptions.	×	0.07
Explicit image corruptions g(x) are found by minimizing the loss function L(xc) using gradient descent.	×	0.03

References

- http://arxiv.org/abs/2306.11113v2
- http://arxiv.org/abs/2206.02435v2
- http://arxiv.org/abs/2107.02926v2