SOVEREIGN: Do node-based BNNs with latent node variables maintain higher accuracy on CIFAR-10-C than weight-based BNNs wh

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 29, 2026

Abstract

Neural Architecture Search (NAS) has shown great success in automating the design of neural networks, but the prohibitive amount of computations behind current NAS methods requires further investigations in improving the sample efficiency and the network evaluation cost to get better results in a shorter time. In this paper, we present a novel scalable Monte Carlo Tree Search (MCTS) based NAS agent, named AlphaX, to tackle these two aspects. AlphaX improves the search efficiency by adaptively balancing the exploration and exploitation at the state level, and by a Meta-Deep Neural Network (DNN)

1 Introduction

Analysis of: Neural Architecture Search Using Deep Neural Networks and Monte Carlo Tree Search. Research goal: Do node-based BNNs with latent node variables maintain higher accuracy on CIFAR-10-C than weight-based BNNs when the number of Monte Carlo samples is reduced to 5?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score $<$ 6.5.

3 Results

9 papers retrieved. 6 claims extracted, 6 verified. Tribunal: 8.3/10 → APPROVE (revision_round=0). Policy: AUTO_APPROVE.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv Relevance ranking is query-dependent. Tribunal consensus is LLM-based and prompt-sensitive.

5 Extracted Claims

Claim	Verified	Confidence
AlphaX found an architecture that reaches 97.84% top-1 accuracy on CIFAR-10	✓	0.21
AlphaX found an architecture that reaches 75.5% top-1 accuracy on ImageNet	✓	0.17
AlphaX is 3x and 2.8x more sample efficient than Random Search and Regularized Evolution in finding the global optimum o	✓	0.28
AlphaX improves the search efficiency by adaptively balancing the exploration and exploitation at the state level	✓	0.26
AlphaX accelerates MCTS rollouts with a distributed design	✓	0.20
AlphaX reduces the number of epochs in evaluating a network by transfer learning	✓	0.20

References

- https://doi.org/10.1007/s10915-022-01939-z
- https://doi.org/10.1609/aaai.v34i06.6554
- https://doi.org/10.1007/s10462-023-10562-9