# SenseNet: Neural Architecture Search Inspired by Adaptive Biological Sensing for Transparency and Adaptability

**Neha Bhargava, Kirtana Sunil Phatnani**
(neha.iitb@gmail.com, kirtana.phatnani@fractal.ai)
Fractal Analytics
Mumbai, Maharashtra, India

## Abstract

Neural Architecture Search (NAS) has emerged as a key method for automating neural network design. However, most existing approaches rely on static optimization strategies that struggle to adapt to new tasks, incorporate real-time feedback, or explain their decision processes—limitations that hinder performance in dynamic environments. Inspired by evolutionary phenomena and human learning — where organisms develop or adapt sensory systems to interpret and act on environmental cues—this work explores how NAS can incorporate similar mechanisms to improve adaptability and dynamic search strategy. This paper presents a hybrid approach called SenseNet that combines large language model (LLM)-driven explainability, dynamic optimization, and real-time adaptability to enhance NAS decision-making. At its core, SenseNet features a meta-controller that dynamically selects high-level strategy - exploration, exploitation, or balanced — based on sensing environmental cues, while an ML strategist translates these high-level strategies into tailored crossover and mutation operations that guide architecture evolution effectively. We evaluate SenseNet comprehensively on NATS-Bench and TransNAS-Bench, demonstrating its adaptability and effectiveness. Our experiments show that SenseNet achieves state-of-the-art results on NATS-Bench and performs competitively on TransNAS-Bench. By embedding sensing and response mechanisms into NAS, SenseNet enhances both efficiency and transparency in neural network optimization, shifting the paradigm from rigid search techniques to biologically inspired, self-adaptive, and transparent neural network optimization.

**Keywords:** Neural Architecture Search (NAS); adaptive optimization; evolutionary computing; real-time feedback; biologically inspired AI; Large Language Models (LLMs); explainable AI (XAI); self-improving AI systems

## Introduction and Motivation

*"It is not the strongest of the species that survives, nor the most intelligent, but the one most responsive to change."*

*— Charles Darwin*

Neural Architecture Search (NAS), a framework for automating the discovery of high-performing neural networks, has demonstrated remarkable success across fields like computer vision, natural language processing, and healthcare. We introduce **SenseNet**, a hybrid NAS framework that integrates the reasoning capabilities of Large Language Models (LLMs) with evolutionary search strategies informed by historical and real-time data. As every species over the course of evolution develops its own *umwelt* — a unique way of sensing and acting in the world — SenseNet gives NAS the ability to sense what's working via strategically using LLMs. Unlike prior LLM-based NAS methods such as GENIUS, GPT-NAS, and EvoPrompting (Zheng et al., 2023; Chen et al., 2024; Rahman & Chakraborty, 2024), which rely on static prompting and discard valuable historical feedback, SenseNet dynamically adapts, interprets, and justifies its choices. Drawing inspiration from cognitive neuroscience and the adaptive mechanisms found in evolution(Sejnowski, 2018), SenseNet advances a new class of NAS systems capable of learning from memory, reasoning across time, and scaling to complex, multi-objective tasks in the real world(White et al., 2023; Chitty-Venkata & Somani, 2022).

## SenseNet Framework

To operationalize our hybrid NAS paradigm, we introduce **SenseNet**, a framework that integrates LLM-driven reasoning with structured, strategy and machine-learning niche-aware evolution. SenseNet is composed of four key modules that collectively emulate the adaptability of a human data scientist: *Meta-Controller* that dynamically selects high-level search strategies (explore, exploit, or balanced) using LLM prompts informed by real-time feedback. *ML Strategist* that translates these strategies into niche-specific architectural adaptations, tailoring mutation and crossover operations to each structural cluster. *QD Optimization with Dual-Niche Framework* Quality-Diversity (QD) optimization framework that preserves performance and diversity both in behavioral strategies and architectural traits, thereby avoiding premature convergence. *LLM-Guided Evolutionary Operations* that apply context-aware crossover and adaptive mutation based on historical trends and niche performance.

Together, these components enable SenseNet to treat NAS not as a fixed algorithmic pipeline, but as a dynamic, evolving perceptual system guided by feedback and interpretability—drawing inspiration from the biological concept of the *Umwelt* (Yong, 2022). By evolving its architectural vision through LLM-mediated feedback loops, SenseNet not only improves search efficiency and solution diversity, but also enhances transparency by maintaining a traceable rationale for every structural decision. This enables it to scale effectively to real-world, multi-objective problems where adaptability and interpretability are paramount.

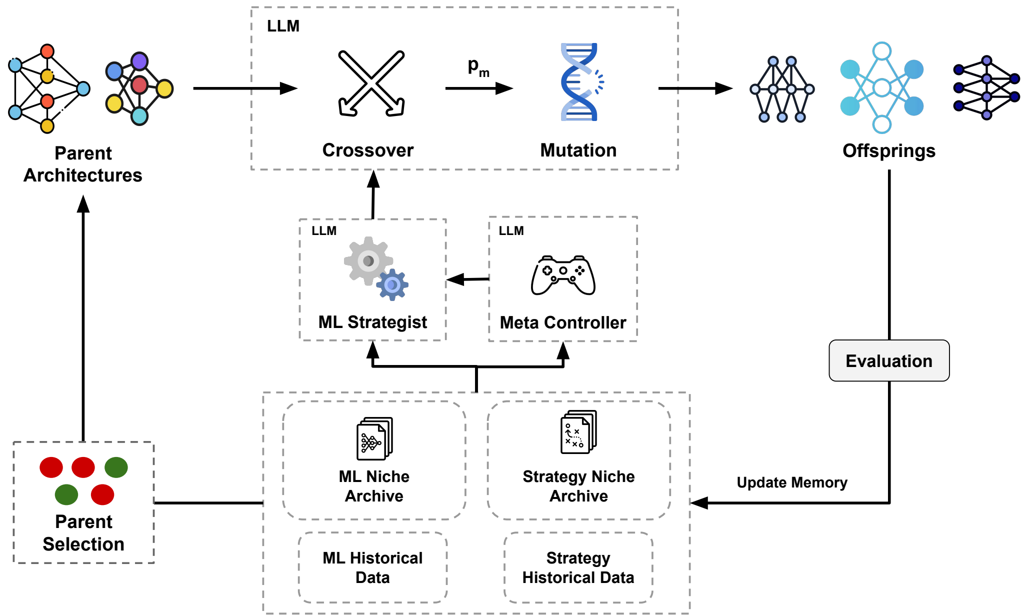


Figure 1: Overview of SenseNet framework.

| Method | CIFAR-10 | | CIFAR-100 | | ImageNet16-120 | |
|---|---|---|---|---|---|---|
| | Val Acc | Test Acc | Val Acc | Test Acc | Val Acc | Test Acc |
| LLMATIC (Nasir, Earle, Togelius, James, & Cleghorn, 2024) | - | 94.26 $\pm$ 0.13 | - | 71.62 $\pm$ 1.73 | - | 45.87 $\pm$ 0.96 |
| EvoPrompting (Chen, Dohan, & So, 2024) | 90.38 $\pm$ 0.33 | 93.11 $\pm$ 0.90 | 70.47 $\pm$ 0.23 | 70.39 $\pm$ 0.58 | 45.32 $\pm$ 0.26 | 45.15 $\pm$ 0.51 |
| $\Lambda$ DARTS (Movahedi et al., 2022) | **91.55** $\pm$ 0.00 | 94.36$\pm$ 0.00 | 73.49 $\pm$ 0.00 | **73.51** $\pm$ 0.00 | 46.37 $\pm$ 0.00 | 46.34 $\pm$ 0.00 |
| IS-DARTS (He, Liu, Zhang, & Zheng, 2024) | **91.55** $\pm$ 0.00 | 94.36$\pm$ 0.00 | 73.49 $\pm$ 0.00 | **73.51** $\pm$ 0.00 | 46.37 $\pm$ 0.00 | 46.34 $\pm$ 0.00 |
| **SenseNet** | 91.41 $\pm$ 0.18 | **94.37** $\pm$ **0.06** | **73.58** $\pm$ **0. 00** | **73.51** $\pm$ **0. 00** | **46.78** $\pm$ **0.33** | **46.5** $\pm$ **0.00** |

Table 1: Performance comparison of SenseNet and various NAS methods on NATS-Bench across CIFAR-10, CIFAR-100, and ImageNet16-120 datasets (Chrabaszcz et al., 2017). SenseNet demonstrates competitive validation and test accuracy (in %).

## Experiments and Results

We evaluated **SenseNet** on NATS-Bench (Mehta et al., 2022; Dong, Liu, Musial, & Gabrys, 2021) and TransNAS-Bench (Duan et al., 2021), leveraging precomputed architecture performance data across diverse datasets including CIFAR-10(Alex, 2009), CIFAR-100, and ImageNet16 (Deng et al., 2009). SenseNet achieved state-of-the-art performance on NATS-Bench (Table 1) and competitive results on TransNAS-Bench , with significantly faster convergence compared to baseline NAS methods.

To understand SenseNet's inner workings, we conducted ablation studies. Removing dynamic strategy selection led to erratic performance trends, while random evolution strategies caused instability—underscoring the value of the Meta-Controller in guiding exploration and exploitation. Eliminating historical data integration resulted in slower convergence and more fluctuation, emphasizing its role in stabilizing decision-making. Together, these experiments validate SenseNet's core components and highlight how hybrid reasoning, feedback integration, and strategy adaptation enhance both efficiency and transparency in neural architecture search.

## Conclusion and future work

This paper introduced **SenseNet**, a hybrid NAS framework that integrates LLM-based reasoning with structured, data-driven strategies to enhance adaptability, transparency, and performance in neural architecture search. While SenseNet demonstrates strong results across NAS benchmarks and offers interpretable, dynamic decision-making, its reliance on LLMs introduces challenges such as potential bias, hallucinations, and computational overhead (Beigi et al., 2024; Dai et al., 2024; Liu et al., 2024; Xu, Jain, & Kankanhalli, 2024). To address these limitations, future work will explore hybrid decision mechanisms that combine LLMs with rule-based logic, expand the diversity of ML niches, and test SenseNet in real-world, open-ended environments. These directions aim to improve robustness, reduce cost, and validate the system's generalizability under less structured conditions. Additionally, incorporating biologically inspired memory mechanisms—such as synaptic plasticity or ambiguity-driven adaptation—may further strengthen SenseNet's capacity for long-term learning and context-aware evolution, advancing the vision of NAS as a cognitively inspired, interpretable, and adaptive system.

## Acknowledgment

We thank Fractal Analytics for their support and for enabling collaboration between the AI Research and CerebrAI teams. We are especially grateful to Aabha Hattangadi, Jieya Rawal, and Tanya Upadhyay for their insightful guidance and meaningful contributions throughout this work.

## References

Alex, K. (2009). Learning multiple layers of features from tiny images. *https://www. cs. toronto. edu/kriz/learning-features-2009-TR. pdf*.

Beigi, M., Wang, S., Shen, Y., Lin, Z., Kulkarni, A., He, J., ... others (2024). Rethinking the uncertainty: A critical review and analysis in the era of large language models. *arXiv preprint arXiv:2410.20199*.

Chen, A., Dohan, D., & So, D. (2024). Evoprompting: language models for code-level neural architecture search. *Advances in Neural Information Processing Systems*, *36*.

Chitty-Venkata, K. T., & Somani, A. K. (2022). Neural architecture search survey: A hardware perspective. *ACM Computing Surveys*, *55*(4), 1–36.

Chrabaszcz, P., Loshchilov, I., & Hutter, F. (2017). A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*.

Dai, S., Xu, C., Xu, S., Pang, L., Dong, Z., & Xu, J. (2024). Bias and unfairness in information retrieval systems: New challenges in the llm era. In *Proceedings of the 30th acm sigkdd conference on knowledge discovery and data mining* (pp. 6437–6447).

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 ieee conference on computer vision and pattern recognition* (pp. 248–255).

Dong, X., Liu, L., Musial, K., & Gabrys, B. (2021). Nats-bench: Benchmarking nas algorithms for architecture topology and size. *IEEE transactions on pattern analysis and machine intelligence*, *44*(7), 3634–3646.

Duan, Y., Chen, X., Xu, H., Chen, Z., Liang, X., Zhang, T., & Li, Z. (2021). Transnas-bench-101: Improving transferability and generalizability of cross-task neural architecture search. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition* (pp. 5251–5260).

He, H., Liu, L., Zhang, H., & Zheng, N. (2024). Is-darts: Stabilizing darts through precise measurement on candidate importance. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 38, pp. 12367–12375).

Liu, F., Liu, Y., Shi, L., Huang, H., Wang, R., Yang, Z., ... Ma, Y. (2024). Exploring and evaluating hallucinations in llm-powered code generation. *arXiv preprint arXiv:2404.00971*.

Mehta, Y., White, C., Zela, A., Krishnakumar, A., Zabergja, G., Moradian, S., ... Hutter, F. (2022). Nas-bench-suite: Nas evaluation is (now) surprisingly easy. *arXiv preprint arXiv:2201.13396*.

Movahedi, S., Adabinejad, M., Imani, A., Keshavarz, A., Dehghani, M., Shakery, A., & Araabi, B. N. (2022). λ-darts: Mitigating performance collapse by harmonizing operation selection among cells. *arXiv preprint arXiv:2210.07998*.

Nasir, M. U., Earle, S., Togelius, J., James, S., & Cleghorn, C. (2024). Llmatic: neural architecture search via large language models and quality diversity optimization. In *Proceedings of the genetic and evolutionary computation conference* (pp. 1110–1118).

Rahman, M. H., & Chakraborty, P. (2024). Lemo-nade: Multi-parameter neural architecture discovery with llms. *arXiv preprint arXiv:2402.18443*.

Sejnowski, T. J. (2018). *The deep learning revolution*. MIT press.

White, C., Safari, M., Sukthanker, R., Ru, B., Elsken, T., Zela, A., ... Hutter, F. (2023). Neural architecture search: Insights from 1000 papers. *arXiv preprint arXiv:2301.08727*.

Xu, Z., Jain, S., & Kankanhalli, M. (2024). Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*.

Yong, E. (2022). *An immense world: How animal senses reveal the hidden realms around us*. Knopf Canada.

Zheng, M., Su, X., You, S., Wang, F., Qian, C., Xu, C., & Albanie, S. (2023). Can gpt-4 perform neural architecture search? *arXiv preprint arXiv:2304.10970*.