Enhanced NAS-PINNs: Strategies for Efficient PDE Solving

Atul Rajput
Dr. B. R. Ambedkar National Institute of Technology Jalandhar
Punjab, India
atulr.mc.24@nitj.ac.in

Neha Yadav*
Dr. B. R. Ambedkar National Institute of Technology Jalandhar
Punjab, India
yadavn@nitj.ac.in

Abstract

In this paper, we propose an improved Neural Architecture Search-guided Physics-Informed Neural Network (PINN) framework named as Enhanced NAS-PINNs for efficiently solving partial differential equations (PDEs). The proposed method integrates differentiable architecture search with adaptive loss balancing, residual-based collocation sampling, and a hybrid optimization strategy combining AdamW and L-BFGS. Further training is stabilized by the addition of exponential moving average (EMA) parameters. Experiments on Poisson, Burgers, and Helmholtz equations demonstrate faster convergence, improved accuracy, and better generalization compared to the baseline NAS-PINN.

1 Introduction

Physics-Informed Neural Networks (PINNs) incorporate physical laws into the loss function, enabling data-efficient solutions of PDEs [Raissi et al., 2019]. Neural Architecture Search (NAS) extends this by automatically discovering optimal architectures [Liu et al., 2019]. Although NAS-PINNs eliminate the need for manual tuning, their efficacy is frequently constrained by optimizer sensitivity, poor selection of collocation sites, and unstable gradient dynamics.

To overcome these challenges, we propose the Enhanced NAS-PINNs algorithm, which combines the search for differentiable architectures with adaptive optimization and sampling strategies. The framework aims to produce architectures that not only fit physics constraints but also train efficiently and robustly across PDE types.

2 Methodology

We formulate NAS-PINNs as a bi-level optimization problem:

$$\min_{\alpha} \ \mathrm{MSE}(\theta^*(\alpha), \alpha), \quad \text{s.t. } \theta^*(\alpha) = \arg\min_{\theta} L(\theta, \alpha),$$

where α are architecture parameters and θ are network weights. The adaptive loss function balances PDE, boundary, and initial condition components:

$$L_{\text{adaptive}} = \sum_{i \in \{\text{PDE,BC,IC}\}} \frac{1}{2\sigma_i^2} L_i + \log \sigma_i.$$

Trainable uncertainty weights σ_i improve gradient stability [Chen et al., 2023]. Residual-based sampling focuses learning on regions with high residuals, updating a fraction of collocation points every K iterations based on $R(x,t) = |F(u_\theta(x,t))|$. A two-stage optimization—AdamW for fast exploration, followed by L-BFGS for refinement—accelerates convergence. An Exponential Moving Average (EMA) of model weights further stabilizes training and improves predictive smoothness.

2nd Conference on Applied AI and Scientific Machine Learning (CASML 2025).

3 Experiments

Experiments were implemented in PyTorch and run on a T4 GPU provided by Google Colab. The NAS process used the differentiable architecture search approach (DARTS) [Liu et al., 2019], with candidate operations such as fully connected layers, various activation functions (tanh, ReLU, sine, GELU) and residual connections. Training used a two-stage optimization strategy:

- AdamW with cosine learning rate decay for fast convergence [Loshchilov and Hutter, 2017].
- L-BFGS for fine-tuning to stabilize residuals.

An Exponential Moving Average (EMA) with $\beta = 0.999$ was applied to enhance stability. Adaptive loss weighting balanced the terms of PDE, boundary, and initial condition, while residual-based adaptive sampling resampled 30% of the collocation points every 1000 iterations. The models were trained for up to 50,000 iterations with early stopping based on validation loss. Accuracy was measured using the relative L2 error and analyzed using convergence counts and error visualizations.

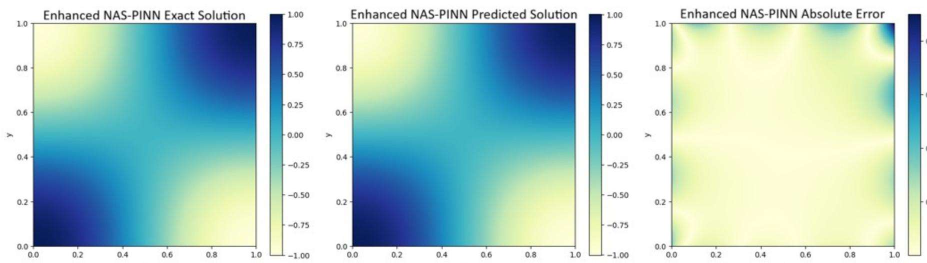

Figure 1: Poisson equation — (a) exact solution, (b) predicted solution, and (c) absolute error distribution.

Enhanced NAS-PINN converge more efficiently than the baseline NAS-PINN and show improved boundary accuracy and phase stability. Adaptive loss weighting maintains balance among PDE, boundary, and initial condition terms, while residual sampling ensures efficient use of collocation points.

4 Conclusion

The paper presents an Enhanced NAS-PINN framework that improves the stability, accuracy, and efficiency of solving PDEs. It incorporates adaptive loss weighting, adaptive sampling, a hybrid optimization strategy AdamW and L-BFGS, and EMA stabilization. These enhancements address issues such as loss imbalance and inefficient sampling. Experiments on the Poisson, Burgers' and Helmholtz equations show faster convergence, higher precision, and better stability compared to the baseline NAS-PINN. The framework also generalizes well to both regular and irregular geometries, with improved gradient stability using the Log-Cosh loss.

5 Acknowledgements

This research was supported by the seed grant (GP-75) from Dr. B. R. Ambedkar National Institute of Technology Jalandhar to the corresponding author and a fellowship from the Ministry of Education, Government of India to the first author.

References

X. Chen, W. Sun, and J. Zhang. Adaptive sampling and loss balancing for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 409:115986, 2023.

H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations (ICLR)*, 2019.

I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.