

# An Efficient Neural Architecture Search Model for Medical Image Classification

Lunchen Xie<sup>1,2</sup>, Eugenio Lomurno<sup>2</sup>, Matteo Gambella<sup>2</sup>, Danilo Ardagna<sup>2</sup>,  
Manuel Roveri<sup>2</sup>, Matteo Matteucci<sup>2</sup>, and Qingjiang Shi<sup>1</sup> \*

<sup>1</sup>Tongji University, Shanghai, China,  
School of Software Engineering

<sup>2</sup>Politecnico di Milano, Milan, Italy,  
Department of Electronics, Information, and Bioengineering

**Abstract.** Accurate classification of medical images is essential for modern diagnostics. Deep learning advancements led clinicians to increasingly use sophisticated models to make faster and more accurate decisions, sometimes replacing human judgment. However, model development is costly and repetitive. Neural Architecture Search (NAS) provides solutions by automating the design of deep learning architectures. This paper presents ZO-DARTS+, a differentiable NAS algorithm that improves search efficiency through a novel method of generating sparse probabilities by bi-level optimization. Experiments on five public medical datasets show that ZO-DARTS+ matches the accuracy of state-of-the-art solutions while reducing search times by up to three times.

## 1 Introduction

Accompanying the advancement of deep learning models, their application in medical image processing has become critically important in contemporary diagnostic processes. In fact, this technology provides experts with the ability to detect diseases earlier and with greater accuracy than previously possible. Although this revolution spans multiple data types and tasks, the design of optimal models remains costly. While image processing models are well established, achieving peak performance requires careful selection of appropriate structures and operations. In addition, the significant heterogeneity and variable quality of medical images necessitate domain-specific adaptation of these models, which requires significant human effort and expertise in neural network design.

To alleviate the burden of repetitive and demanding tasks from humans, the Neural Architecture Search (NAS) paradigm has been proposed and is nowadays considered a silver bullet for automating the model design process [1]. As this approach is relatively new to the field of medical diagnosis, it holds considerable potential for application in medical image processing, particularly in classification tasks. Moreover, given the complexity of the initial NAS techniques, novel trends concerning efficient search strategies and hardware-aware techniques are currently being developed to enhance the performance and accessibility of deep learning solutions [2, 3, 4].

---

\*This work was supported in part by the National Key R&D Program of China under grant 2022YFA1003900, in part by the NSFC under Grant 62231019, in part by the FAIR project funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence), and in part by the China Scholarship Council (CSC).

The main contribution of this paper is the introduction of ZO-DARTS+, an efficient and accurate differentiable NAS algorithm tailored for medical image classification. This new algorithm extends ZO-DARTS [5] by incorporating sparsemax [6] along with an appropriate annealing strategy. These enhancements allow the algorithm to generate operation probabilities with sparse values during the search, improving interpretability. In addition, ZO-DARTS+ converges faster than its predecessor, reducing the average search time by 26.6% without sacrificing, and even improving, the accuracy of the resulting architectures. When compared with other state-of-the-art solutions, ZO-DARTS+ achieves comparable levels of accuracy, while reducing the search time by up to three times.

## 2 Background

The cell-based approach is prevalent in NAS to reduce search space as it can shrink the search space from model-level to cell-level. Inside each cell, operations between every  $i^{\text{th}}$  and  $j^{\text{th}}$  data block utilise a shared candidate set  $\mathcal{O} = \{o^{(i,j)}\}$ . DARTS introduces architecture parameters  $\alpha = \{\alpha^{(i,j)}\}$ , which are continuous and form a weighted sum of the operations [7]. Each  $\alpha^{(i,j)} \in \mathbb{R}^{|\mathcal{O}|}$  acts as the weight vector for the operation set  $\{o^{(i,j)}\}$  and has to be optimized. The mixed operation is normalized by the softmax function and defined as follows:

$$\bar{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(x). \quad (1)$$

This introduces a bi-level optimization problem with  $\alpha$  as the upper-level variable and model parameters  $\mathbf{w}$  as the lower-level variable:

$$\min_{\alpha} F(\alpha) = \mathcal{L}_{val}(\mathbf{w}^*(\alpha), \alpha) \quad s.t. \quad \mathbf{w}^*(\alpha) = \operatorname{argmin}_{\mathbf{w}} \mathcal{L}_{train}(\mathbf{w}, \alpha). \quad (2)$$

Here,  $\mathcal{L}_{val}(\mathbf{w}^*, \alpha)$  and  $\mathcal{L}_{train}(\mathbf{w}, \alpha)$  are loss functions on the validation and training datasets respectively. The final discrete architecture is derived by applying argmax over  $\{\alpha^{(i,j)}\}$ .

To solve Eq. (2) accurately, the analytical gradient of  $F(\alpha)$  is derived using the implicit function theorem [8]:

$$\nabla_{\alpha} F(\alpha) = \nabla_{\alpha} \mathcal{L}_{val}(\mathbf{w}^*, \alpha) + \nabla_{\alpha}^T \mathbf{w}^*(\alpha) \nabla_{\mathbf{w}} \mathcal{L}_{val}(\mathbf{w}^*(\alpha), \alpha), \quad (3)$$

where

$$\nabla_{\alpha} \mathbf{w}^*(\alpha) = -[\nabla_{\mathbf{w}\mathbf{w}}^2 \mathcal{L}_{train}(\mathbf{w}^*, \alpha)]^{-1} \nabla_{\alpha\mathbf{w}}^2 \mathcal{L}_{train}(\mathbf{w}^*, \alpha). \quad (4)$$

Due to the impracticality of computing the Hessian as in Eq. (4), Xie et al. introduce ZO-DARTS [5] to circumvent this obstacle. ZO-DARTS employs a zeroth-order approximation technique [9]:

$$\tilde{\nabla}_{\alpha} F(\alpha) \triangleq \frac{F(\alpha + \mu \mathbf{u}) - F(\alpha)}{\mu} \mathbf{u}. \quad (5)$$

As  $\mu \rightarrow 0$ ,  $\frac{F(\alpha + \mu \mathbf{u}) - F(\alpha)}{\mu}$  approximates the directional derivative  $\nabla_{\alpha}^T F(\alpha) \mathbf{u}$ .

Combining Eq. (3) and Eq. (5) we have:

$$\begin{aligned} \tilde{\nabla}_{\alpha} F(\alpha) &\approx \nabla_{\alpha}^T F(\alpha) \mathbf{u} \mathbf{u} = \nabla_{\alpha}^T \mathcal{L}_{val}(\mathbf{w}^*, \alpha) \mathbf{u} \mathbf{u} + \nabla_{\mathbf{w}}^T \mathcal{L}_{val}(\mathbf{w}^*(\alpha), \alpha) \nabla_{\alpha} \mathbf{w}^*(\alpha) \mathbf{u} \mathbf{u} \\ &= \nabla_{\alpha}^T \mathcal{L}_{val}(\mathbf{w}^*, \alpha) \mathbf{u} \mathbf{u} + [\nabla_{\alpha} \mathbf{w}^*(\alpha) \mathbf{u}]^T \nabla_{\mathbf{w}} \mathcal{L}_{val}(\mathbf{w}^*(\alpha), \alpha) \mathbf{u}. \end{aligned} \quad (6)$$

Here  $\nabla_{\alpha} \mathbf{w}^*(\alpha) \mathbf{u}$  can be substituted by  $\frac{\mathbf{w}^*(\alpha + \mu \mathbf{u}) - \mathbf{w}^*(\alpha)}{\mu}$  according to the definition of directional derivative, too. Thus, Eq. (6) can be computed swiftly.

### 3 The proposed method

ZO-DARTS+ is an advanced NAS algorithm that improves efficiency by integrating a novel probability normalization function with an annealing strategy. Traditionally, softmax is used to generate probability weights from a set of architecture parameters, which are then used to compute a weighted sum of operations in DARTS as stated in Section 2. However, this method struggles with the complexity of bi-level NAS problems, as softmax tends to produce probabilities that do not sufficiently diverge and remain too similar across different operations. This similarity is detrimental to the selection of discrete operations, where a sparse probability distribution is preferable. To overcome these limitations, ZO-DARTS+ employs the sparsemax function [6], which projects the input vector  $\alpha$  onto the probability simplex to produce a sparser probability distribution:

$$\text{sparsemax}(\alpha) := \operatorname{argmin}_{\mathbf{p} \in \Delta^{K-1}} \|\mathbf{p} - \alpha\|^2. \quad (7)$$

The  $(K - 1)$ -dimensional simplex,  $\Delta^{K-1} := \{\mathbf{p} \in \mathbb{R}^K | \mathbf{1}^T \mathbf{p} = 1, \mathbf{p} \geq 0\}$ , ensures that the sum of the probabilities is equal to one and each probability is non-negative, increasing the likelihood of obtaining a sparse result. This sparse expression of the probabilities is a proper normalization function to substitute softmax as used in Eq. (1):

$$\bar{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \text{sparsemax}_o(\alpha) o(x). \quad (8)$$

Furthermore, sparsemax is designed to push architecture parameters to extremes (close to 1 for selected operations and 0 for others), but it does not always eliminate non-zero elements completely. For softmax, a temperature parameter  $\tau$  is commonly used to modulate the input values thus increasing sparsity:

$$\text{softmax}'_o(\alpha) = \frac{\exp(\alpha_o^{(i,j)}/\tau)}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)}/\tau)}. \quad (9)$$

As  $\tau \rightarrow 0$  this approach closely resembles an argmax function, ideal for achieving discrete selections. We effectively integrate this property into sparsemax by combining the trick of dividing input  $\alpha_o^{(i,j)}$  by  $\tau$ . To further enhance this property, we propose an annealing strategy to reduce  $\tau$  by a manually chosen factor  $a \leq 1$  every  $m$  epochs, driving the system towards a one-hot vector. Therefore, the final normalization function used in ZO-DARTS+ is:

$$\text{sparsemax}'(\alpha) = \operatorname{argmin}_{\mathbf{p} \in \Delta^{K-1}} \|\mathbf{p} - \alpha / (\tau * a^{n//m})\|^2, \quad (10)$$

where  $n$  is the total number of epochs, and  $//$  indicates an integer division. By adjusting the factors  $a$  and  $m$ , the algorithm efficiently accelerates the search process while promoting sparsity, closely matching the need for fast and automatic model generation.

### 4 Experiments and Results

To demonstrate the effectiveness and efficiency of our approach, we conducted experiments on five medical image datasets from MedMNIST [10], namely PneumoniaMNIST, OCTMNIST, BreastMNIST, OrganAMNIST, and OrganSMNIST.

Table 1: Accuracy comparison with baselines and DARTS-style methods.

Dataset	PneumoniaMNIST	OCTMNIST	BreastMNIST	OrganAMNIST	OrganSMNIST
ResNet18	0.854	0.743	0.863	<b>0.935</b>	0.782
AutoKeras	0.878	0.763	0.831	0.905	<b>0.813</b>
Google AutoML	<b>0.946</b>	0.771	0.861	0.886	0.749
DARTS	0.935±0.008	0.756±0.028	0.876±0.010	<b>0.935±0.014</b>	0.771±0.015
MiLeNAS	0.625±0.000	0.610±0.312	<b>0.885±0.000</b>	0.695±0.442	0.762±0.008
ZO-DARTS	<u>0.941±0.003</u>	<u>0.807±0.005</u>	0.878±0.013	<u>0.928±0.011</u>	<u>0.787±0.004</u>
ZO-DARTS+ (ours)	0.939±0.008	<b>0.815±0.019</b>	<b>0.891±0.006</b>	0.925±0.018	0.779±0.008

Table 2: Search Time (second) comparison with DARTS-style methods.

Dataset	PneumoniaMNIST	OCTMNIST	BreastMNIST	OrganAMNIST	OrganSMNIST
DARTS	1990.3	4217.1	503.9	2173.8	754.2
MiLeNAS	1726.4	3436.1	390.7	1895.1	644.0
ZO-DARTS	825.5	1792.3	213.7	874.8	354.6
ZO-DARTS+ (ours)	<b>605.1</b>	<b>1381.4</b>	<b>150.7</b>	<b>629.0</b>	<b>263.9</b>

In our NAS framework, both ZO-DARTS and ZO-DARTS+ update the architecture parameters  $\alpha$  only after the model parameters  $\mathbf{w}$  are optimized 10 times. An initial temperature  $\tau = 1.5$  is used to encourage exploration in the early stages of the search. The annealing factor is set to  $a = 0.75$ , with an interval of  $m = 5$ , allowing for a gradual refinement of the model configurations. For benchmarking, we used ResNet18 as a standard for manually designed models and included AutoKeras and Google AutoML Vision as baselines for NAS models. In addition, we evaluated three DARTS-style methods – DARTS [7], MileNAS [11], and ZO-DARTS [5] – all of which operate within a search space defined by NAS-Bench-201 [12] with five operations: Zeroise, Skip Connect, 1x1 Conv, 3x3 Conv, and 3x3 Avg Pooling. We performed 50 search rounds (epochs) for each method. The final model accuracy is determined by retraining to full convergence after each search. We conducted experiments three times to get the average and standard deviation of performance, and the average search time.

Table 1 shows the optimal performance achieved by each algorithm on different datasets, with baseline data taken directly from the MedMNIST homepage<sup>1</sup>. ZO-DARTS+ outperforms ResNet-18, Google AutoML, and AutoKeras on three datasets while having a close performance gap for the rest. It also consistently ranks first or close to the best among DARTS-style methods, demonstrating robust performance across the board. Despite this narrow accuracy gap between ZO-DARTS+ and other DARTS-style methods, there is a significant reduction in search time. Surprisingly, ZO-DARTS+ tends to converge and stop the search by early stopping around the 40<sup>th</sup> epoch of search, further reducing the search time. As detailed in Table 2, this efficiency underlines the suitability of ZO-DARTS+ for the development of medical image classification models, balancing state-of-the-art accuracy with up to a threefold reduction in search time.

We analyzed the progression of architecture operator parameters across the different NAS methods and plotted these progressions in Fig 1. All models started optimization with the same initial probability weights. Throughout the

<sup>1</sup><https://medmnist.com>

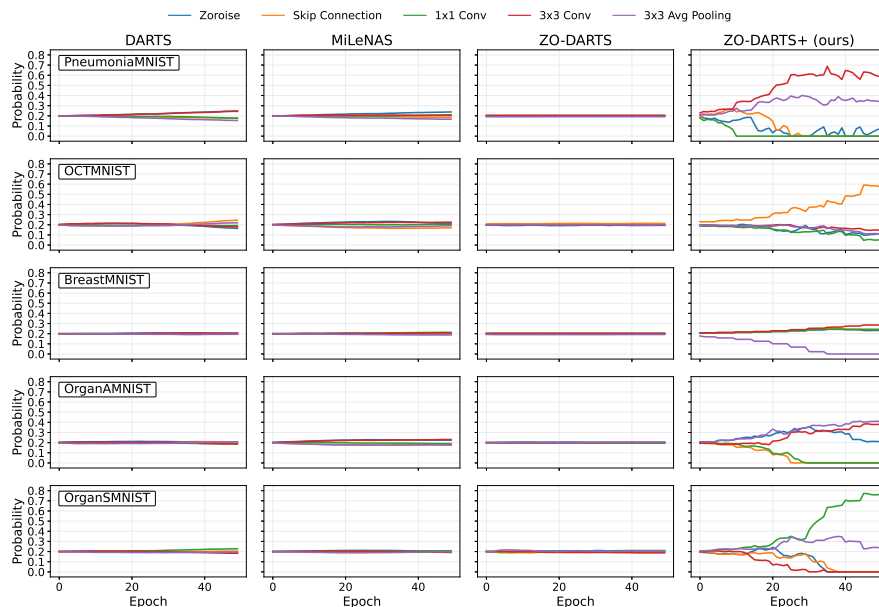


Fig. 1: Probability rank variation of one edge during the search procedure.

search process, most NAS methods exhibited limited variability in these weights, leading to challenges in reliably selecting the most effective operation – the probabilities remained too close together, allowing for easy rank changes, thus compromising the interpretability of the operators. Unlike other methods, ZO-DARTS+ showed rapid convergence in selecting the best operation, thanks to the integration of sparsemax and an annealing strategy. This efficiency not only improves performance, but also avoids the high costs associated with increasing the number of search rounds to achieve clear probability separations.

In addition, the analysis revealed distinct preferences for certain operations of one edge depending on the dataset. For example, the 3x3 Conv operation was frequently selected for the PneumoniaMNIST dataset, whereas the 3x3 Avg Pooling was predominant in the OCTMNIST dataset, with other operations often excluded from the final selection altogether. This pattern highlights the adaptability of ZO-DARTS+, allowing it to effectively tailor its architecture to different types of data. These findings – coupled with robust accuracy results and reduced search times – confirm that ZO-DARTS+ is a highly validated model that demonstrates significant adaptability and reliability in its operation selection, making it a superior choice for medical image classification tasks.

## 5 Conclusions and Future works

This paper presents ZO-DARTS+, a novel differentiable NAS algorithm that achieves up to three times faster search times compared to existing leading methods while preserving accuracy on medical image datasets. These results validate our hypothesis and are achieved through the integration of the sparsemax function and a tailored annealing strategy during operator selection, resulting in

sparser and more efficient solutions. Future research will build on this foundation by applying stringent constraints to reduce both computational resources and time further. In addition, the application of sparsity-aware annealing strategies in different NAS frameworks will be explored to validate the effectiveness of the approach in different settings.

## References

- [1] Matt Poyser and Toby P. Breckon. Neural architecture search: A contemporary literature review for computer vision applications. *Pattern Recognition*, 147:110052, 2024.
- [2] Andrea Falanti, Eugenio Lomurno, Danilo Ardagna, and Matteo Matteucci. Popnasv3: a pareto-optimal neural architecture search solution for image and time series classification. *Applied Soft Computing*, 145:110555, 2023.
- [3] Matteo Gambella, Alessandro Falcetta, and Manuel Roveri. CNAS: Constrained Neural Architecture Search. In *SMC*. IEEE, 2022.
- [4] Jeon-Seong Kang, JinKyu Kang, Jung-Jun Kim, Kwang-Woo Jeon, Hyun-Joon Chung, and Byung-Hoon Park. Neural architecture search survey: A computer vision perspective. *Sensors*, 23(3):1713, 2023.
- [5] Lunchen Xie, Kaiyu Huang, Fan Xu, and Qingjiang Shi. ZO-DARTS: Differentiable architecture search with zeroth-order approximation. In *ICASSP*. IEEE, 2023.
- [6] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *ICML*. PMLR, 2016.
- [7] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. In *ICLR*, 2019.
- [8] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *AISTATS*, pages 1540–1552, 2020.
- [9] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *FoCM*, 17(2):527–566, 2017.
- [10] Jiancheng Yang, Rui Shi, Donglai Wei, Zequan Liu, Lin Zhao, Bilian Ke, Hanspeter Pfister, and Bingbing Ni. Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data*, 2023.
- [11] Chaoyang He, Haishan Ye, Li Shen, and Tong Zhang. Milenas: Efficient neural architecture search via mixed-level reformulation. In *CVPR*, 2020.
- [12] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *ICLR*, 2020.