Interpreting Operation Selection in Differentiable Architecture Search: A Perspective from Influence-Directed Explanations

Anonymous Author(s)
Affiliation
Address
email

Abstract

The *Differentiable ARchiTecture Search* (DARTS) has dominated the neural architecture search community due to its search efficiency and simplicity. DARTS leverages continuous relaxation to convert the intractable operation selection problem into a continuous magnitude optimization problem which can be easily handled with gradient-descent, while it poses an additional challenge in measuring the operation importance or selecting an architecture from the optimized magnitudes. The vanilla DARTS assumes the optimized magnitudes reflect the importance of operations, while more recent works find this naive assumption leads to poor generalization and is without any theoretical guarantees. In this work, we leverage influence functions, the functional derivatives of the loss function, to theoretically reveal the operation selection part in DARTS and estimate the candidate operation importance by approximating its influence on the supernet with Taylor expansions. We show the operation strength is not only related to the magnitude but also second-order information, leading to a fundamentally new criterion for operation selection in DARTS, named ***Influential Magnitude***. Empirical studies across different tasks on several spaces show that vanilla DARTS and its variants can avoid most failures by leveraging the proposed theory-driven operation selection criterion.

1 Introduction

Neural Architecture Search (NAS) has been successfully applied to automating the process of neural network design for a broad of deep learning fields [35]. However, the early NAS methods are often heavily computational-expensive and require hundreds or even thousands GPU days [34]. To improve the efficiency for more practical applications, a lot of recent works are proposed based on one-shot (also known as weight-sharing) paradigm [33] to improve the search efficiency significantly. *Differentiable ARchiTecture Search* (DARTS) [27], as the most popular one-shot NAS method, has dominated this area since its appearance. DARTS leverages continuous relaxation to convert the intractable operation selection problem into a continuous magnitude optimization problem, which can be gracefully solved by gradient-descent based on the bi-level optimization [8]. A discrete architecture is then derived based on the optimized magnitudes through operation selection.

Although the continuous relaxation makes DARTS able to optimize the magnitudes with rigorous theoretical formulation and guarantees [49] in the supernet training part, it poses an additional challenge in selecting a discrete architecture from the optimized magnitudes that there is no existing work can give a theoretical explanation. DARTS is based on the weight-sharing paradigm [33], which assumes that the derived discrete architecture with the trained supernet weights can predict its train-from-scratch performance. However, there is a gap between the two groups of weights, as depicted in Figure 1 in Appendix. The vanilla DARTS assumes the optimized magnitudes can

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

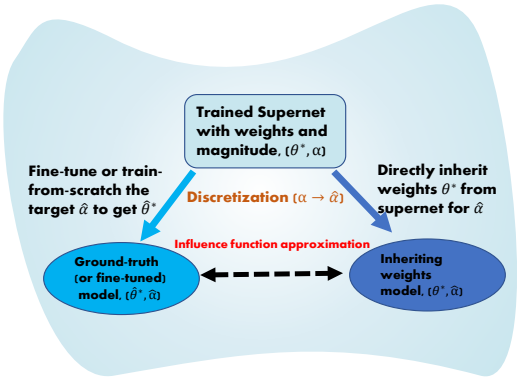

Figure 1: **Pictorial depiction of discretization part in DARTS.** There is a gap between the train-from-the-scratch (or fine-tuned) weights $\hat{\theta}^*$ based on $\hat{\alpha}$ after discretization and the trained supernet weights θ^*. This paper leverages influence functions to quantify the disparity between $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$ and $\mathcal{L}(\theta^*, \hat{\alpha})$ and show how weights θ changes with α. The discretization process can be the *argmax* adopted by DARTS [27], the perturbation-based DARTS-PT [40], and so on.

reflect the importance of the candidate operations, while Wang et al. [40] empirically show that the operation magnitude does not necessarily indicate how important the operation is in a trained supernet. For example, DARTS prefers *skip-connection* [45] as its magnitude always dominates the remaining operations, while an architecture with intensive *skip-connections* generally leads to poor performance. Accordingly, Wang et al. [40] proposed a heuristic perturbation-based method, DARTS-PT, to measure each candidate operation's influence on the supernet by alternatively removing each one to calculate the supernet performance drop. Despite DARTS-PT can partially relieve the failure of DARTS, we empirically found that DARTS-PT still prefers *skip-connection* than other operations. One potential reason is that *skip-connection*'s magnitude is much higher than other candidate operations, and directly removing *skip-connections* usually greatly deteriorates the supernet performance. In addition, DARTS-PT lacks theoretical explanations and guarantees. For example, the intuitive way to show the importance of a candidate operation importance is the leave-one-out retraining [22], while which should fine-tune the supernet to get the performance drop after removing every operation rather than directly measuring the performance drop as DARTS-PT, since there also exists a gap between the two groups of weights, as visualized in Figure 1. Section 3 and Appendix C also give an analysis on why perturbation-based method fails.

This paper focuses on filling the gap of theoretical explanations for the discrete operation selection part in DARTS. Instead of considering a heuristic approach to measure the operation importance [27, 40], this paper first rigorously shows how the architecture parameters discretization poses an influence on the supernet weights in DARTS, and estimates the supernet performance drop accordingly. Our method is motivated by influence function [22], which is originally designed to understand the effect of removing an individual training point, while we leverage Taylor's approximation to interpret the supernet performance drop when removing a candidate operation. In this way, we show the operation strength is not only related to the magnitude but also second-order information, and we accordingly design a new criterion for the operation selection in DARTS, called ***Influential Magnitude*** ($\mathcal{I}_{\mathcal{M}}$). Then, to avoid calculating the inverse of the Hessian matrix with respect to the model weights, we further utilize the Neumann series [29] and Sherman-Morrison formula [37] to approximate the Inverse-Hessian-Vector Products (IHVPs), making the proposed method applicable to DARTS. Finally, we verify the effectiveness of our method on several NAS benchmark datasets and the common DARTS search space. To summarize, we make the following contributions:

- This paper deepens our understanding of the operation selection in DARTS. We reformulate the operation selection in DARTS by approximating its influence on the supernet with Taylor expansions, which can thus gracefully interpret how the validation performance changes when selecting different operations without any additional fine-tuning.
- We theoretically reveal the operation strength is not only related to the magnitude but also the second-order information, and accordingly derive a fundamentally new criterion to measure the operation sensitivity, which we call ***Influential Magnitude*** ($\mathcal{I}_{\mathcal{M}}$), for architecture

selection in DARTS. Further, to make the proposed criterion practical, we devise several methods to estimate the Inverse-Hessian-Vector Products (IHVPs) in calculating the second-order information.

- Extensive experiments verify the effectiveness of the proposed criterion, which significantly improves the performance of vanilla DARTS and other baselines on the NAS benchmark datasets and the common DARTS search space. We show DARTS is still a strong NAS baseline when considering a more reasonable and theory-driven operation strength metric.

2 Preliminaries: DARTS and Influence Function

Typically, NAS focuses on finding a cell structure that is represented as a directed acyclic graph (DAG) with *N* nodes and *E* edges affiliated with operations, which is a discrete and hard-to-optimize problem. To enable the gradient-descent for architecture search, **DARTS** [27] leverages ***softmax*** as the continuous relaxation function to transform the architecture search into a continuous magnitude optimization problem:

$$\alpha_o = \frac{\exp(\bar{\alpha}_o)}{\sum_{o' \in \mathcal{O}} \exp(\bar{\alpha}_{o'})},$$

where $\mathcal{O}$ contains all candidate operations and DARTS is to optimize the magnitude α_o, which can be formulated as a bi-level optimization problem [8]:

$$\begin{aligned} &\min_{\alpha} \quad \mathcal{L}_{val}(\theta^*(\alpha), \alpha) \\ &\text{s.t.} \quad \theta^*(\alpha) = \text{argmin}_{\theta}\ \mathcal{L}_{train}(\theta, \alpha), \end{aligned} \tag{1}$$

where α is the continuous architecture representation and θ is the supernet weights. The nested formulation in DARTS is the same as the gradient-based hyperparameter optimization with bi-level optimization, which has been validated in a broad applications [13, 31, 32]. After the bi-level optimization, DARTS usually considers heuristic methods to derive the final architecture, where the most popular one is to select the operations with the highest magnitudes, $\hat{\alpha} = \boldsymbol{argmax}(\alpha)$. Despite its simplicity, more recent works [7, 40, 45, 47, 48] found that DARTS could hardly find satisfactory architectures.

Different from the majority of existing works that attribute the failure of DARTS to its magnitude optimization part, which is under a rigorous theoretical formulation [49], we revisit DARTS from the perspective of the architecture selection and try to reveal this part with theoretical explanations, where we aim to identify the operation that will contribute most to the supernet performance, through the lens of influence functions [22] to mathematically model the architecture selection process instead of heuristics.

The influence function [18] is a classic technique from robust statistics that reveals how model parameters change as we upweight or perturb a specific training sample, which has been applied in explaining a lot of modern machine learning applications [2, 22, 30, 50]. More specifically, given a model with parameters θ^* after trained on a full training set $\mathcal{D}$, influence function aims to study the change of parameters $\hat{\theta}^* - \theta^*$ when retraining the model parameters to $\hat{\theta}^*$ after deleting a specific training point z, that $\hat{\mathcal{D}} = \mathcal{D} - z$. Rather than retraining the model for every removed point, Koh and Liang [22] estimated the influence by upweighting the data point with some small ϵ. With leveraging the Taylor expansions, the influence on the model parameters can be estimated as:

$$\mathcal{I}(z, \theta) = \frac{d\hat{\theta}^*}{d\epsilon} = -H_{\theta^*}^{-1} \nabla_\theta \mathcal{L}(z, \theta^*), \tag{2}$$

and its influence on the loss function of a test point z_{test} is accordingly calculated as:

$$\mathcal{I}(z, \mathcal{L}) = \frac{d\mathcal{L}(z_{\text{test}}, \hat{\theta}^*)}{d\epsilon} = -\nabla_\theta \mathcal{L}(z_{\text{test}}, \theta^*)^\top H_{\theta^*}^{-1} \nabla_\theta \mathcal{L}(z, \theta^*). \tag{3}$$

The detailed derivation of the two influence functions can be found in [22]. Based on Eq.(2) and (3), we can now approximate the validation performance drop when removing a training point z from the training set without conducting the time-consuming retraining.

The original work of influence functions for modern machine learning [22], which is also called as first-order influence functions, only considers a very small change on θ^* as it only changes one

training point, so the first-order Taylor expansion is satisfied to calculate the influence. However, upweighting a group of data points may lead to a large change on θ^* [4, 5], which violates the small perturbation assumption of the first-order influence functions. To make the the influence function capture the model parameter change more precisely, Basu et al. [4] extended influence functions using second-order approximations under the assumption that the third-order derivatives of the loss function at optimum is small. For the quadratic loss, the third-order derivatives of the loss are zero or very small, and Basu et al. [4] empirically verified that the this assumption also approximately held for the classification problem with cross-entropy loss function.

Different from existing works that analyze the effects on model parameters of removing data points, this paper creatively adapts influence functions to estimate the importance of candidate operations on a trained supernet in differentiable architecture search. We also follow the assumption in [4] that neglect the third-order and higher derivatives of the loss, and analyze the operation selection part, which is an additional process after bi-level supernet training in DARTS, through the lens of influence function in the following sections.

3 Interpret Operation Selection through Influence Functions

This section aims to understand how the model parameters θ (a.k.a. the supernet weights in DARTS) would change when we derive the discrete architecture from the optimized continuous operation magnitudes, a.k.a. architecture selection. In this paper, our analysis only focuses on this architecture selection part after the bi-level supernet training process is finished. To be more clearly, we first define the optimized operation magnitudes by DARTS as α with the optimized supernet weights θ^*, and the validation loss is $\mathcal{L}(\theta^*, \alpha)$ for simplicity. Then, we define the architecture parameters change to $\hat{\alpha}$ after discretization, and the corresponding ground-truth supernet weights and validation loss are defined as $\hat{\theta}^*$ and $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$, respectively. Rather than deleting a single data point that only brings small changes on the model parameters [22], the change of architecture parameters can lead to a considerable change on θ, so we leverage the second-order approximation [4] to reveal the supernet weights change. With second-order Taylor expansion on $\hat{\theta}^*$ for $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$, we have:

$$\Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \approx \mathcal{L}(\theta^*, \hat{\alpha}) + \Delta\theta^T \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} + 1/2 \Delta\theta^T \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta} \Delta\theta - \mathcal{L}(\theta^*, \alpha), \tag{4}$$

where $\Delta\theta = \hat{\theta}^* - \theta^*$, and $\mathcal{L}$ is the validation loss function in Eq.(1). For example, when we only remove one candidate operation as DARTS-PT [40] to measure the operation importance, the influence of this discretization brings to θ is $\Delta\theta$, and the true validation loss change should be $\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha)$ as stated in Eq.(4) rather than $\mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha)$ used by DARTS-PT. As shown in Eq.(4), we need to calculate $\Delta\theta$ to get $\Delta\mathcal{L}$ after we change the α to $\hat{\alpha}$, while we try to avoid the time-consuming retraining process through leveraging influence functions [22] to estimating $\hat{\theta}^*$. Based on the implicit function theorem [29, 49], the validation loss change after discretizing α to $\hat{\alpha}$ can be approximated as the following theorem.

Theorem 1 *Suppose that DARTS obtains the optimized architecture parameter α with supernet weights θ^* after supernet training, α changes to $\hat{\alpha}$ when conducting architecture discretization, and the train-from-scratch validation loss for $\hat{\alpha}$ is $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$. If the third and higher derivatives of the loss function $\mathcal{L}$ at optimum is zero or sufficiently small [4], and with $\frac{\partial \mathcal{L}(\hat{\theta}^*, \hat{\alpha})}{\partial \theta} = 0$, we have*

$$\Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \approx \mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) - 1/2 \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta}^{-1} \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}. \tag{5}$$

Theorem 1 states that we can estimate the train-from-scratch validation performance of $\hat{\alpha}$ without any retraining. During the discretization process, we can apply the ***argmax*** on all edges by one-shot that $\hat{\alpha} = \boldsymbol{argmax}(\alpha)$, or only discretize one edge in each step, or remove one candidate operation from one edge as DARTS-PT [40] which can be formulated as the first part of Eq.(5), $\mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha)$. We prefer an iterative discretization process that only applies a small change on α as DARTS-PT, since our theoretical analysis in the next corollary will show that the error bound of our approximation is related to $\Delta\alpha = \hat{\alpha} - \alpha$. Before our analysis, we give the following common assumptions, which are also considered in several recent bi-level optimization problems [9, 15–17, 49].

Assumption 1 *For any θ and α, $\mathcal{L}(\cdot,\alpha)$ and $\mathcal{L}(\theta,\cdot)$ are Lipschitz continuous with $C_f > 0$ and $C_L > 0$, respectively.*

Assumption 2 *$\mathcal{L}(\theta,\alpha)$ is twice differentiable with constant C_H and is λ-strongly convex with θ around $\theta^*(\alpha)$.*

Assumption 3 *$\left\|\nabla^2_{\theta\alpha}\mathcal{L}\right\|$ is bounded with constant $C_a > 0$.*

Corollary 1 *Based on the Assumption 1-3, we could bound the error between the approximated validation loss $\mathcal{L}(\hat{\theta}^*,\hat{\alpha}) = \Delta\mathcal{L} + \mathcal{L}(\theta^*,\alpha)$ and the ground-truth $\tilde{\mathcal{L}}(\hat{\theta}^*,\hat{\alpha})$ in DARTS with $E = \left\|\mathcal{L}(\hat{\theta}^*,\hat{\alpha}) - \tilde{\mathcal{L}}(\hat{\theta}^*,\hat{\alpha})\right\| \leqslant \frac{K^3}{6}\max\left|\frac{\partial\mathcal{L}^3}{\partial\theta^3}\right|$, where $K = \frac{C_L}{\lambda}\left\|\Delta\alpha\right\| + \frac{C_H C_a^2}{2\sigma^2_{min}\lambda}\left\|\Delta\alpha\right\|^2 + o(\left\|\Delta\alpha\right\|^2)$.*

Compared with the *argmax* used in DARTS, we prefer the perturbation-based discretization method [41] as which only remove one operation at each step, since the above Corollary 1 show that error bound of the approximated validation loss $\mathcal{L}(\hat{\theta}^*,\hat{\alpha})$ increase with the magnitude of the change of α.

4 Influential Magnitude for Operation Selection in DARTS

Section 3 provides an iterative solution to measure the operation importance based Eq. (5) through individually removing each operation, while which needs to repeat n times (the number of all candidate operations for all edges in the supernet work). More important, since removing one operation poses a considerable change on α ($\alpha_i \to 0$), the approximation error produced by Eq.(5) is non-negligible which may affects the accuracy as stated by **Corollary** 1. So, a more practical solution to illustrate the importance of each operation is to estimate how the validation performance will change after posing an infinitesimal change on α, a.k.a. operation sensitivity. In this section, we aim to investigate the operation importance through the lens of sensitivity [24]. Different from Section 3 that removes one operation ($\alpha_i \to 0$), this section only changes the α_i with an infinitesimal ε_i, which means that we can now conduct first-order Taylor expansion on α for $\mathcal{L}(\hat{\theta}^*,\hat{\alpha})$. Further with second-order Taylor expansion on θ and assuming the local optimal $\frac{\partial\mathcal{L}(\theta^*,\alpha)}{\partial\theta} = 0$, we have:

$$\Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*,\hat{\alpha}) - \mathcal{L}(\theta^*,\alpha) \approx 1/2\Delta\theta^T\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\theta}\Delta\theta + \Delta\alpha^T\frac{\partial\mathcal{L}(\hat{\theta}^*,\alpha)}{\partial\alpha}. \tag{6}$$

Similar to Eq.(4), it is intractable to directly calculate $\Delta\theta$, while we can leverage implicit function theorem to approximate it. Theorem 2 presents how the validation performance changes after we pose an infinitesimal change on α.

Theorem 2 *Suppose that DARTS obtains the optimized architecture parameter α with supernet weights θ^* after supernet training, and we pose an infinitesimal change on α. Based on implicit function theorem and under the assumption that the third and higher derivatives of the loss function at optimum is zero or sufficiently small [4], the change of validation loss can be estimated as:*

$$\Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*,\hat{\alpha}) - \mathcal{L}(\theta^*,\alpha) \approx -1/2\Delta\alpha^T\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}H^{-1}\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha} * \Delta\alpha, \tag{7}$$

where $H = \frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\theta}$ is the Hessian matrix.

From Theorem 2, we can observe the relationship between $\Delta\mathcal{L}$ and $\Delta\alpha$. To be more specifically, when we only consider an infinitesimal change on α, that $\Delta\alpha = \varepsilon\cdot\mathbf{1}$ where $\mathbf{1}$ is a column vector with all ones and ε is an infinitesimal scalar, the sensitivity of α can be defined as:

$$\frac{\Delta\mathcal{L}}{\Delta\alpha} = -1/2\cdot\varepsilon\cdot\mathbf{1}^T\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}H^{-1}\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha}. \tag{8}$$

Accordingly, we propose an ***Influential Magnitude*** to quantify the operation importance through the lens of sensitivity:

Definition 1 ***Influential Magnitude*** *($\mathcal{I}_\mathcal{M}$): Suppose DARTS obtains the optimized magnitude α with supernet weights θ^* after supernet training, the operation sensitivity can be defined as $\mathcal{I}_\mathcal{M} = -\boldsymbol{1}^T\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}H^{-1}\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha}$.*

Definition 1 shows we can replace the magnitude α with the operation sensitivity $\mathcal{I}_\mathcal{M}$ in the operation selection for DARTS, whose *i-th* value represents the importance of *i-th* candidate operation.

5 Practical Calculation on Operation Influence

For a large neural network, it is impractical to calculate the second-order information, e.g., the Hessian matrix H, let alone the inverse of Hessian. Generally, the core challenge in calculation of Eq.(5) and Eq.(7) is the Inverse-Hessian Vector Products (IHVPs), $H^{-1}v$ [14, 22, 37]. As to the term of $\frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha}$, we follow DARTS [27] to utilize the finite difference approximation for the practical calculation. The first technique for IHVPs is the conjugate gradient (CG) which converts the matrix inverse problem into an optimization problem with δ-optimal solution. However, the standard CG can be slow with a large dataset as it needs to go through the whole training points in each step [1]. The original paper [22] using influence function for modern machine learning problem leverages the Neumann series to approximate the inverse, that $H^{-1} = \sum_{k=0}^{\infty}(I-H)^k$, with only consider the first t terms to approximate H^{-1}. However, there is a strict condition in the Neumann series approximation that $\|I - H\| \leq 1$, which can not be guaranteed in the practical implementation. Differently, we consider that $H^{-1} = \gamma(\gamma H)^{-1} = \gamma\sum_{k=0}^{\infty}(I-\gamma H)^k$, where γ is a small enough value. The following lemma gives our Neumann series based implementation for IHVPs.

Lemma 1 *With small enough γ, and assuming $\mathcal{L}$ is λ-strongly convex at optimum, $H^{-1}v$ can be formulated as: $H^{-1}v = \gamma\sum_{k=0}^{K}[I-\gamma H]^k V_0 = V_0 + V_1 + ... + V_K$, where $H = \frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\theta}$, $V_0 = \gamma v$, and $V_1 = \gamma(I-\gamma H)V_0, ..., V_K = \gamma(I-\gamma H)V_{K-1}$.*

In the practical implementation, it is still hard to make sure $\mathcal{L}$ is λ-strongly convex and not easy to select an appropriate γ. The work [5] empirically showed that the stochastic Neumann series estimation for IHVP is somehow erroneous, especially when the network is deep. As we know, the true Fisher information matrix converges to the Hessian matrix with the training loss approaches to zero. However, it is usually infeasible to calculate the exact Hessian or Fisher matrix in a large neural network. A practical approximation is called as empirical Fisher, which has been verified to be effective in variety of applications [3, 19]. In addition, an important advantage from empirical Fisher is that it allows to approximate the inverse of the Hessian based on the Sherman-Morrison formula. Although the empirical Fisher can not exactly match the true Fisher and there are several active researches on its applicability [23, 38], empirical Fisher is significantly more computationally-efficient, which obtains a good trade-off between the approximation and practical efficiency [37]. Apart from Neumann series, we also consider Sherman-Morrison formula [14, 37] $(A + \mathrm{u}\mathrm{v}^T)^{-1} = A^{-1} - \frac{A^{-1}\mathrm{u}\mathrm{v}^T A^{-1}}{1+\mathrm{v}^T A^{-1}\mathrm{u}}$ to calculate IHVPs. With empirical Fisher $F = \sum_{j=1}^{n}\nabla_\theta\mathcal{L}_j(\nabla_\theta\mathcal{L}_j)^T$ replacing true Fisher $\hat{F} = \nabla_\theta f(\nabla_\theta f)^T$, IHVPs can be calculated as following.

Lemma 2 *When assume the empirical Fisher can approximate the Fisher matrix, and H is the Hessian matrix $\frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\theta}$ in the optimal point, the IHVPs $H^{-1}v$ can be formulated as: $H^{-1}v = F_n^{-1}v = F_{n-1}^{-1}v - \mathrm{r}_n\frac{\mathrm{r}_n{}^T v}{N+\nabla_\theta\mathcal{L}_n{}^T\mathrm{r}_n} = \eta^{-1}v - \sum_{j=1}^{n}\mathrm{r}_j\frac{\mathrm{r}_j{}^T v}{N+\nabla_\theta\mathcal{L}_j{}^T\mathrm{r}_j}$, where $\mathcal{L} = \ell + \eta\mathcal{R}(\theta)$ that ℓ is a cross-entropy loss and $\mathcal{R}$ is the regularization term, $F_n = \frac{1}{n}\sum_{j=1}^{n}\nabla_\theta\mathcal{L}_j\nabla_\theta\mathcal{L}_j^T$ is the empirical Fisher, and $\mathrm{r}_j = F_{j-1}^{-1}\nabla_\theta\mathcal{L}_j$ which can be recurrently calculated through $\mathrm{r}_j = \eta^{-1}\nabla_\theta\mathcal{L}_j - \sum_{i=1}^{j-1}\mathrm{r}_i\frac{\mathrm{r}_i{}^T\nabla_\theta\mathcal{L}_j}{N+\nabla_\theta\mathcal{L}_i{}^T\mathrm{r}_i}$.*

6 Experiments

In Section 3 and 4, we have theoretically interpreted the operation selection through individually measuring each operation's influence on supernet, and accordingly devised an ***Influential Magnitude*** ($\mathcal{I}_\mathcal{M}$) to replace the commonly-used magnitude α for architecture selection. In this section, we conduct a series of experiments to analyze whether our proposed method leads to better and more stable results in differentiable architecture search settings. We consider two NAS benchmark dataset search space, NAS-Bench-201 [12] and NAS-Bench-1shot1 [46], and the most popular DARTS search space [27]. We verified the effectiveness of the proposed theory-driven method compared with two heuristic methods, *argmax* and

Table 1: Test accuracy (%) on NAS-Bench-201.

Method	CIFAR-10	CIFAR-100	ImageNet-16-120
Random	86.61±13.46	60.83±12.58	33.13±9.66
DARTS	54.30±0.00	15.61±0.00	16.32±0.00
DARTS-PT	89.63±0.19	62.35±2.14	36.51±2.13
DARTS-IF	91.84±0.84	67.94±1.23	42.50±3.30
DARTS-IM	**93.61±0.23**	**71.31±0.40**	**44.98±0.36**

Full comparison with SOTAs is in Appendix D. We default use Sherman-Morrison for DARTS-IM in this experiment.

perturbation-based selection, and show that the proposed ***Influential Magnitude*** can be applied to DARTS and its variants with greatly enhancing the performance.

6.1 Reproducible Comparison on Benchmark Datasets

To empirically verify the effectiveness of our method, we first run a series of experiments on the two NAS benchmark datasets, NAS-Bench-201 and NAS-Bench-1shot1. To make our results reproducible, we only replace the selection part in DARTS with $\mathcal{I}_{\mathcal{M}}$ and keep the remaining part identical.

Table 1 summarizes the statistical comparison results on NAS-Bench-201 with different *random seeds*, where DARTS-IF, as described in Section 3, follows DARTS-PT that individually removes operation to measure the supernet performance drop, while DARTS-IM, as described in Section 4, calculates the operation sensitivity (influential magnitude) to replace the original magnitude. Algorithm 1 and 2 in the Appendix H sketched the two proposed methods. The **Random** baseline is to randomly select architectures without considering α or θ, and DARTS and DARTS-PT leverage the *argmax* and perturbation-based heuristic selection, respectively. As shown, our DARTS-IF and DARTS-IM both significantly outperform the random baseline and the two heuristic baselines, where DARTS-IM achieves near-optimal results, a **94.29**%, **72.67**%, and **45.93**% test accuracy on CIFAR-10, CIFAR-100, and ImageNet, respectively, in a single run (DARTS-IM with Sherman-Morrison approximation and $N = 20$ under *random seed* 0).

Figure 8 in Appendix G plots the searched architectures with different architecture selection methods in NAS-Bench-201. We can see that DARTS selects *skip-connection* for all edges as its magnitude usually outweighs other candidate operations. The results of DARTS in NAS-Bench-201, which are even poorer than the random baseline, also suggest that the magnitudes could hardly reflect the operation importance. Although DARTS-PT can partially relieve this issue by selecting one or two more convolutional operations, it still prefers *skip-connections* (selects 3 *skip-connections* for 6 edges) where directly removing them still significantly affect the performance. On the contrary, our DARTS-IF and DARTS-IM can find valid architectures with more convolutional operations (with only one *skip-connection*). Results in Table 1 verified the effectiveness of our influence functions based selection methods, which outperform baselines by large margins. DARTS-IF is similar to DARTS-PT while with an additional term to approximate the fine-tuned performance and the results that DARTS-IF outperforms DARTS-PT verified the effectiveness of the influence function explanation.

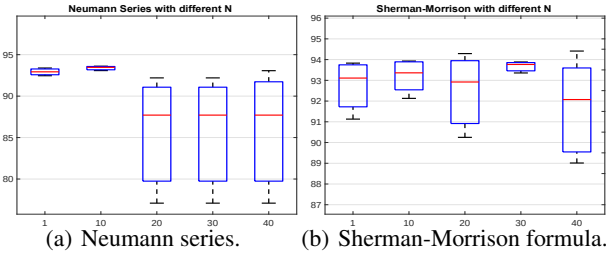

(a) Neumann series. (b) Sherman-Morrison formula.

Figure 2: Ablation study on N under two approximation methods, where x-axis is N and y-axis represents test accuracy on CIFAR-10. The computational complexity increases with the square of N for Sherman-Morrison formula, and we set the maximal $N = 40$.

As discussed in Section 5, we can leverage different methods to approximate the IHVPs in our $\mathcal{I}_{\mathcal{M}}$, and Figure 2 also analyzes different approximation methods along with the number of batches N. Full results on three datasets can be found in the Figure 5 in Appendix F. As demonstrated, increasing N can help the Sherman-Morrison approximation to more accurately catch the operation importance and we empirically find that $N = 30$ is enough in the NAS-Bench-201 dataset. However, we also observe Neumann approximation is not stable which even obtains worse results with increasing number of batches, showing it is somehow erroneous in estimating IHVPs. Appendix F analyzes different inverse Hessian estimation methods for our $\mathcal{I}_{\mathcal{M}}$.

Table 2: Best test error (%) on NAS-Bench-1shot1.

Method	Space1	Space2	Space3
DARTS	6.17±0.09	6.30±0.00	6.80±0.00
DARTS-PT	6.25±0.05	**6.28±0.06**	6.69±0.21
DARTS-IM	**6.10±0.24**	6.53±0.05	**6.20±0.00**
PC-DARTS	6.37±0.05	6.30±0.00	6.50±0.00
PC-DARTS-PT	6.14±0.08	6.37±0.12	6.38±0.09
PC-DARTS-IM	**5.90±0.24**	**6.20±0.22**	**6.10±0.08**

The NAS-Bench-1shot1 is another popular NAS benchmark dataset to analyze differentiable NAS methods, which is built based on the NAS-Bench-101 [44], through dividing all architectures in

NAS-Bench-101 into three different unified cell-based search spaces to enable differentiable NAS methods to be directly applied. NAS-Bench-1shot1 contains 6240, 29160, and 363648 architectures in three spaces, respectively, where Space3 is much more complicated than the remaining two spaces. In NAS-Bench-1shot1, we tracked the searched architecture in every epoch under different operation selection paradigms. We considered two popular supernet training baselines, DARTS and PC-DARTS, and Table 2 summarised the best-searched architectures during the supernet training in different sub-search spaces in NAS-Bench-1shot1. As demonstrated, our $\mathcal{I}_{\mathcal{M}}$ can enhance the performance of DARTS and its variant PC-DARTS in different search spaces of NAS-Bench-1shot1, showing that our influence magnitude (IM) can find more competitive architectures than the default *argmax*, especially in the more complicated space. Compared with the perturbation strategy (PT), our IM is more effective to measure the operation sensitivity with obtaining 5 better results in 6 cases.

6.2 Reproducible Comparison on DARTS Search Space

The DARTS search space is the most commonly-used space to evaluate the NAS methods, while most SOTAs only report the performance of their best searched architectures without analyzing the search statistics. To encourage reproducibility, we train the supernet with the same random seeds, and the trained supernet is utilized to derive the architecture under different operation selection strategies. The comparison results are provided in Table 3, from which we can see that, compared with the *argmax* baseline, the architectures searched by our ***Influential Magnitude*** obtain better performance on CIFAR-10 and ImageNet than DARTS and its variant PC-DARTS, and DARTS-PT on ImageNet, implying our method is more stable to find valid and competitive architectures. Figure 9 in Appendix G also plots the searched architecture examples with different selection methods in the DARTS space. As shown, although DARTS, PC-DARTS and DARTS-PT can find satisfying architecture after multiple runs, the *argmax* is hard to consistently obtain stable results. In addition, different from the perturbation-based operation selection which needs to fine-tune the supernet after the discretization of each edge that costs 0.4 additional GPU day for DARTS, our influential magnitude is calculated by one time which only costs hundred seconds, making our method is much more efficient. Detailed comparisons on best-searched architecture in DARTS space with the SOTAs can be found in Table 5 in Appendix E.

Table 3: Search results on DARTS space.

Method	CIFAR-10 Test Error (%) Single	Multi*	ImageNet Best
DARTS	2.76±0.09	3.02±0.45	26.9 / 8.7
PC-DARTS	2.57±0.07	2.92±0.26	25.1 / 7.8
DARTS-PT	**2.61±0.08**	2.89±0.31	26.1 / 8.2
DARTS-IM	**2.50±0.10**	**2.70±0.18**	**25.0 / 7.6**

* We run the architecture search with multiple times, and average the different derived architecture's test error.

6.3 Discussion on the Robustness of Search

The robustness has been a critical concern in the differentiable architecture search community since Zela et al. [45] observed a performance collapse that DARTS tends to yield degenerate architectures, which especially prefers the *skip-connections* with the search progressing. Similar to DARTS-PT, we also focus on the operation selection part to explain the failure of DARTS. Figure 3 presents several architectures by DARTS with our $\mathcal{I}_{\mathcal{M}}$ on different search spaces. Apart from NAS-Bench-201, we also consider several tool search spaces (S1-S4) proposed by [45], which are specifically designed to verify the robustness of differentiable NAS methods. The complete analysis and searched architectures can be found in Appendix G. As shown in Figure 3, DARTS-IM can find more meaningful architectures in different spaces. For example, DARTS fails dramatically in most spaces that only selects *skip-connection* for all edges, while

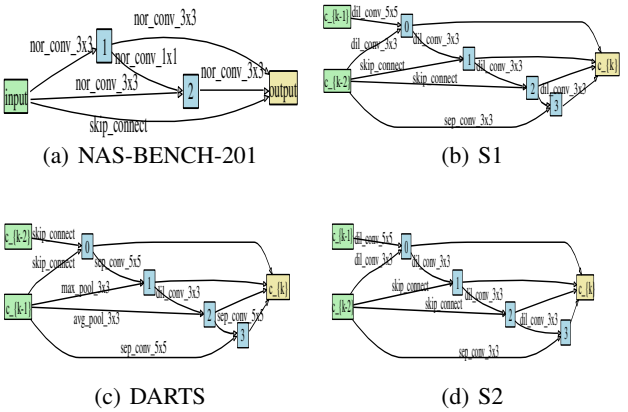

(a) NAS-BENCH-201 (b) S1

(c) DARTS (d) S2

Figure 3: Searched architectures on different search spaces.

DARTS-IM can find architecture with 5 convolutions in NAS-Bench-201, and 6 and 8 convolutions in S1 and S2.

6.4 Discussion at Non-Convergence

As discussed in Section 4, our influential magnitude is derived when θ is in the optimum, with leveraging different methods to approximate the IHVPs. In this subsection, we analyze whether the proposed influential magnitude is still effective to find important operations when θ is not in the optimum. Figure 4 tracked the performance of searched architecture in each epoch during the search with Sherman-Morrison formula. Full results on three datasets and more comparison with DARTS-PT can be found in the Figure 7 in Appendix F. As shown, our influential magnitude $\mathcal{I}_{\mathcal{M}}$ is still effective when θ is not in the optimum, while which is getting more stable with the supernet training. Specifically, as shown in Figure 4 (a) and (b), our $\mathcal{I}_{\mathcal{M}}$ with the different N are both effective, implying that our method is sufficiently informative for the operation selection even under different approximation noises.

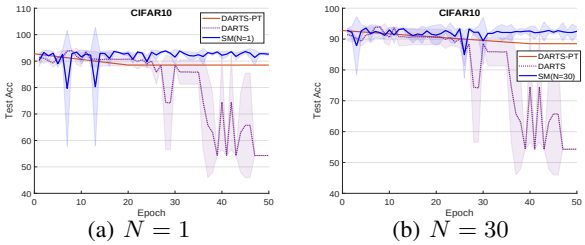

(a) $N = 1$ (b) $N = 30$

Figure 4: Track the architectures during search on NAS-Bench-201 with Sherman-Morrison under different N.

More interesting, our DARTS-IM under Sherman-Morrison formula with N=1 in Figure 4 means that we simply approximate the inverse $H^{-1} = \eta^{-1}\left(I - \eta^{-1}\frac{\nabla_\theta \mathcal{L} \nabla_\theta \mathcal{L}^\top}{1+\eta^{-1}\nabla_\theta \mathcal{L}^\top \nabla_\theta \mathcal{L}}\right) \propto \hat{H}^{-1}$, that $\hat{H}^{-1}$ is a diagonal matrix that $\hat{H}_{ii} = (\nabla_{\theta_i}\mathcal{L}\nabla_{\theta_i}\mathcal{L})$, which is also called the diagonal approximation of the inverse Hessian. As shown, with roughly approximated second-order information, Eq. (7) also help DARTS-IM to find more important operations compared with existing heuristic methods.

6.5 Discussion on the Limitations in Practice

Although our influential magnitude provides a theory-driven and concise solution for the operation selection in DARTS, we consider several assumptions and approximations for the practical implementation. For example, Basu et al. [4, 5] pointed out that the second-order influence values still underestimated the true group influence values in the deep neural networks, while they also observed the correlation marginally improves with the number of parameters in a group and the second-order influence function generally outperforms the first-order. Fortunately, this paper considers the whole parameters in the supernet, which helps the second-order influence function reach its potential.

The approximation of IHVPs requires the conditions that the Hessian matrix is positive definite [29] or the empirical Fisher can replace the true Fisher matrix [14], and the supernet weight θ^* is a global minimum of the empirical risk, while which in practice, are hard to hold. In the calculation of $v * \frac{\partial^2 \mathcal{L}}{\partial\alpha\partial\theta}$, which is detailed described in Appendix B, this paper also follows DARTS that considers the finite difference approximation for practical calculation. These approximations and assumptions may bring errors for the practical calculation, while empirical results verified that our influential magnitude is sufficiently informative for the operation selection in DARTS.

7 Conclusion and Future Work

This paper focuses on the operation selection, which is an essential part while with less attention in the differentiable architecture search. Rather than considering heuristics, this paper introduces the influence functions to theoretically reveal the candidate operation's importance through approximating its influence on the supernet. By leveraging the Taylor's approximation, we gracefully interpret how the validation performance changes during the discretization in DARTS and show that the operation strength is not only related to the magnitude but also the second-order information, where an influential magnitude is accordingly devised for the architecture selection. We based our framework on DARTS and extensive experimental results verified the proposed influential magnitude can be practically used in operation selection to avoid failures in differentiable architecture search. Developing more accurate influential magnitude for architecture search or leveraging the training dynamics [20] to reveal the operation influence in the early of training and also extending influential magnitude to other bi-level optimization problems[21, 28, 29], e.g., network pruning [28], are among directions for future work.

References

[1] N. Agarwal, B. Bullins, and E. Hazan. Second-order stochastic optimization in linear time. *stat*, 1050:15, 2016.

[2] A. Alaa and M. Van Der Schaar. Validating causal inference models via influence functions. In *International Conference on Machine Learning*, pages 191–201. PMLR, 2019.

[3] S.-I. Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.

[4] S. Basu, X. You, and S. Feizi. On second-order group influence functions for black-box predictions. In *International Conference on Machine Learning*, pages 715–724. PMLR, 2020.

[5] S. Basu, P. Pope, and S. Feizi. Influence functions in deep learning are fragile. In *ICLR*, 2021.

[6] F. P. Casale, J. Gordon, and N. Fusi. Probabilistic neural architecture search. *arXiv preprint arXiv:1902.05116*, 2019.

[7] X. Chen and C.-J. Hsieh. Stabilizing differentiable architecture search via perturbation-based regularization. In *ICML*, 2020.

[8] B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of operations research*, 153(1):235–256, 2007.

[9] N. Couellan and W. Wang. On the convergence of stochastic bi-level gradient methods. *Optimization*, 2016.

[10] X. Dong and Y. Yang. Searching for a robust neural architecture in four gpu hours. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 2019.

[11] X. Dong and Y. Yang. One-shot neural architecture search via self-evaluated template network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3681–3690, 2019.

[12] X. Dong and Y. Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. *ICLR*, 2020.

[13] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *International Conference on Machine Learning*, pages 1568–1577. PMLR, 2018.

[14] E. Frantar, E. Kurtic, and D. Alistarh. M-fac: Efficient matrix-free approximations of second-order information. *Advances in Neural Information Processing Systems*, 34, 2021.

[15] S. Ghadimi and M. Wang. Approximation methods for bilevel programming. *arXiv preprint arXiv:1802.02246*, 2018.

[16] R. Grazzi, L. Franceschi, M. Pontil, and S. Salzo. On the iteration complexity of hypergradient computation. In *International Conference on Machine Learning*, pages 3748–3758. PMLR, 2020.

[17] R. Grazzi, M. Pontil, and S. Salzo. Convergence properties of stochastic hypergradients. In *International Conference on Artificial Intelligence and Statistics*, pages 3826–3834. PMLR, 2021.

[18] F. R. Hampel. The influence curve and its role in robust estimation. *Journal of the american statistical association*, 69(346):383–393, 1974.

[19] B. Hassibi and D. Stork. Second order derivatives for network pruning: Optimal brain surgeon. *Advances in neural information processing systems*, 5, 1992.

[20] H. He and W. J. Su. The local elasticity of neural networks. In *ICLR*, 2020.

[21] W. Jin, Y. Ma, X. Liu, X. Tang, S. Wang, and J. Tang. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 66–74, 2020.

[22] P. W. Koh and P. Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pages 1885–1894. PMLR, 2017.

[23] F. Kunstner, P. Hennig, and L. Balles. Limitations of the empirical fisher approximation for natural gradient descent. *Advances in neural information processing systems*, 32, 2019.

[24] N. Lee, T. Ajanthan, and P. Torr. Snip: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2019.

[25] N. Lee, T. Ajanthan, S. Gould, and P. H. Torr. A signal propagation perspective for pruning neural networks at initialization. In *International Conference on Learning Representations*, 2020.

[26] L. Li and A. Talwalkar. Random search and reproducibility for neural architecture search. *arXiv preprint arXiv:1902.07638*, 2019.

[27] H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search. *ICLR*, 2019.

[28] L. Liu, S. Zhang, Z. Kuang, A. Zhou, J.-H. Xue, X. Wang, Y. Chen, W. Yang, Q. Liao, and W. Zhang. Group fisher pruning for practical network compression. In *International Conference on Machine Learning*, pages 7021–7032. PMLR, 2021.

[29] J. Lorraine, P. Vicol, and D. Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1552. PMLR, 2020.

[30] K. Lu, Z. Wang, P. Mardziel, and A. Datta. Influence patterns for explaining information flow in bert. *Advances in Neural Information Processing Systems*, 34, 2021.

[31] D. Maclaurin, D. Duvenaud, and R. Adams. Gradient-based hyperparameter optimization through reversible learning. In *International conference on machine learning*, pages 2113–2122. PMLR, 2015.

[32] F. Pedregosa. Hyperparameter optimization with approximate gradient. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 737–746, 2016.

[33] H. Pham, M. Guan, B. Zoph, Q. Le, and J. Dean. Efficient neural architecture search via parameter sharing. In *International Conference on Machine Learning*, pages 4092–4101, 2018.

[34] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. *AAAI*, 2019.

[35] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen, and X. Wang. A comprehensive survey of neural architecture search: Challenges and solutions. *arXiv preprint arXiv:2006.02903*, 2020.

[36] A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots. Truncated back-propagation for bilevel optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1723–1732. PMLR, 2019.

[37] S. P. Singh and D. Alistarh. Woodfisher: Efficient second-order approximation for neural network compression. *Advances in Neural Information Processing Systems*, 33:18098–18109, 2020.

[38] V. Thomas, F. Pedregosa, B. Merriënboer, P.-A. Manzagol, Y. Bengio, and N. Le Roux. On the interplay between noise and curvature and its effect on optimization and generalization. In *International Conference on Artificial Intelligence and Statistics*, pages 3503–3513. PMLR, 2020.

[39] C. Wang, G. Zhang, and R. Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020.

[40] R. Wang, M. Cheng, X. Chen, X. Tang, and C.-J. Hsieh. Rethinking architecture selection in differentiable nas. In *ICLR*, 2021.

[41] R. Wang, M. Cheng, X. Chen, X. Tang, and C.-J. Hsieh. Rethinking architecture selection in differentiable nas. In *International Conference on Learning Representations*, 2021.

[42] S. Xie, H. Zheng, C. Liu, and L. Lin. Snas: stochastic neural architecture search. *ICLR*, 2019.

[43] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong. Pc-darts: Partial channel connections for memory-efficient architecture search. In *ICLR*, 2020.

[44] C. Ying, A. Klein, E. Christiansen, E. Real, K. Murphy, and F. Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *ICML*, pages 7105–7114, 2019.

[45] A. Zela, T. Elsken, T. Saikia, Y. Marrakchi, T. Brox, and F. Hutter. Understanding and robustifying differentiable architecture search. In *ICLR*, 2020.

[46] A. Zela, J. Siems, and F. Hutter. Nas-bench-1shot1: Benchmarking and dissecting one-shot neural architecture search. In *ICLR*, 2020.

[47] M. Zhang, H. Li, S. Pan, X. Chang, and S. Su. Overcoming multi-model forgetting in one-shot nas with diversity maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7809–7818, 2020.

[48] M. Zhang, H. Li, S. Pan, X. Chang, C. Zhou, Z. Ge, and S. W. Su. One-shot neural architecture search: Maximising diversity to overcome catastrophic forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[49] M. Zhang, S. Su, S. Pan, X. Chang, E. Abbasnejad, and R. Haffari. idarts: Differentiable architecture search with stochastic implicit gradients. In *International Conference on Machine Learning*. PMLR, 2021.

[50] W. Zhang, Y. Wang, Z. You, M. Cao, P. Huang, J. Shan, Z. Yang, and B. Cui. Rim: Reliable influence-based active learning on graphs. *Advances in Neural Information Processing Systems*, 34, 2021.

[51] X. Zheng, R. Ji, L. Tang, B. Zhang, J. Liu, and Q. Tian. Multinomial distribution learning for effective neural architecture search. In *International Conference on Computer Vision (ICCV)*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [No]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proof

Theorem 1: Suppose that DARTS obtains the optimized architecture parameter α with supernet weights θ^* after supernet training, α changes to $\hat{\alpha}$ when conducting architecture discretization, and the train-from-scratch validation loss for $\hat{\alpha}$ is $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$. If the third-derivative of the loss function $\mathcal{L}$ at optimum is zero or sufficiently small [4], and with $\frac{\partial \mathcal{L}(\hat{\theta}^*, \hat{\alpha})}{\partial \theta} = 0$, we have

$$\begin{aligned} \Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \approx \mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \\ - 1/2 * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}. \end{aligned} \tag{9}$$

Proof of Theorem 1: Based on Eq.(4), we now only need to calculate $\Delta\theta$ to get $\Delta\mathcal{L}$ after we change the α to $\hat{\alpha}$ as other parts could be directly obtained, while we try to avoid the time-consuming retraining process through leveraging implicit function to estimating $\hat{\theta}^*$. First, based on the local optimal, we have

$$\frac{\partial \mathcal{L}(\hat{\theta}^*, \hat{\alpha})}{\partial \theta} = 0. \tag{10}$$

When we consider the second-order Taylor expansion on θ, we have

$$\frac{\partial(\mathcal{L}(\theta^*, \hat{\alpha}) + \Delta\theta * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} + \Delta\theta * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta} * \Delta\theta^T)}{\partial \theta} \approx 0. \tag{11}$$

$$\frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} + \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta} * \Delta\theta \approx 0. \tag{12}$$

Eq. (12) is obtained since we assume the third-derivative of the loss function $\mathcal{L}$ at optimum is zero or sufficiently small [4]. And we have

$$\Delta\theta \approx -\frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}. \tag{13}$$

So,

$$\Delta\theta^T \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} \approx -\frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} \tag{14}$$

$$\Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta} * \Delta\theta \approx \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} \tag{15}$$

In this way, based on Eq.(4), (14) and (15), we have

$$\begin{aligned} \Delta\mathcal{L} &= \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \\ &\approx \mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) - 1/2 * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T \\ &* \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} \end{aligned} \tag{16}$$

In this way, we can approximate every change in the α.

□

Corollary 1: Based on the Assumption 1-3, we could bound the error between the approximated validation loss $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$ and the ground-truth $\tilde{\mathcal{L}}_2(\hat{\theta}^*, \hat{\alpha})$ in DARTS with $E = \left\| \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \tilde{\mathcal{L}}(\hat{\theta}^*, \hat{\alpha}) \right\| \leqslant \frac{K^3}{6} \max \left| \frac{\partial \mathcal{L}^3}{\partial \theta^3} \right|$, where $K = \frac{C_L}{\lambda} * \|\Delta\alpha\| + \frac{C_H * C_a^2}{2 * \sigma_{min}^2 * \lambda} * \|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^4)$.

Proof of Corollary 1: Before our analysis on the error bound of approximation in Theorem 1, we first restate the following common assumptions in the bi-level optimization [9, 15–17, 49].

Assumption 1: For any θ and α, $\mathcal{L}(\cdot, \alpha)$ and $\mathcal{L}(\theta, \cdot)$ are Lipschitz continuous with constant $C_f > 0$ and constant $C_L > 0$, respectively.

Assumption 2: $\mathcal{L}(\theta, \alpha)$ is twice differentiable with constant C_H and is λ-strongly convex with θ around $\theta^*(\alpha)$.

Assumption 3: $\left\|\nabla^2_{\theta\alpha}\mathcal{L}\right\|$ is bounded with constant $C_a > 0$.

Now we can estimate the error bound produced by Theorem 1. First, the error is due to the second-order Taylor expansion on θ in Eq. (4). We have the error

$$E = \left\|\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \tilde{\mathcal{L}}(\hat{\theta}^*, \hat{\alpha})\right\| \leq \frac{|\Delta\theta|^3}{6} \max \left|\frac{\partial \mathcal{L}^3}{\partial \theta^3}\right|, \tag{17}$$

where we should notice that the approximated validation loss $\mathcal{L}(\hat{\theta}^*, \hat{\alpha})$ is Taylor expansion of $\hat{\mathcal{L}}$ in the point $\hat{\theta}^* - \Delta\theta$. Since we set that $\Delta\theta = -\frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta}$ in Eq. 13, this error E could be easily bound. However, we should notice that our $\Delta\theta = -\frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta}$ is an approximation, that $\theta^* + \Delta\theta \neq \hat{\theta}^*$, and we define that $\theta^* + \Delta\theta = \hat{\theta}^*_e$ and $\Delta\theta_e = \hat{\theta}^* - \hat{\theta}^*_e$.

First, we have

$$\|\Delta\theta\| = \left\|\theta^* - \hat{\theta}^* + \hat{\theta}^* - \hat{\theta}^*_e\right\| \leq \left\|\theta^* - \hat{\theta}^*\right\| + \left\|\hat{\theta}^* - \hat{\theta}^*_e\right\| \tag{18}$$

For the first part in Eq. (18), based on Assumption 2 that $\mathcal{L}(\theta, \alpha)$ is λ-strongly convex with θ around θ^*, we have:

$$\left\|\nabla\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \nabla\mathcal{L}(\theta^*, \hat{\alpha})\right\| \geq \lambda * \left\|\hat{\theta}^* - \theta^*\right\|, \tag{19}$$

and

$$\left\|\theta^* - \hat{\theta}^*\right\| \leq \frac{1}{\lambda} * \|\nabla\mathcal{L}(\theta^*, \hat{\alpha}^*)\|. \tag{20}$$

When we consider that $\mathcal{L}(\theta^*, \alpha)$ is C_L Lipschitz continuous with α, we have

$$\|\nabla\mathcal{L}(\theta^*, \hat{\alpha}^*) - \nabla\mathcal{L}(\theta^*, \alpha^*)\| \leq C_L * \|\Delta\alpha\|. \tag{21}$$

Baesd on Eq. (20) and (21), we have

$$\left\|\theta^* - \hat{\theta}^*\right\| \leq \frac{C_L}{\lambda} * \|\Delta\alpha\| \tag{22}$$

As to the second part, when we assume that $\mathcal{L}(\theta, \alpha)$ is λ-convex near the local optimal, we also have

$$\left\|\nabla\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \nabla\mathcal{L}(\hat{\theta}^*_e, \hat{\alpha})\right\| \geq \lambda * \left\|\hat{\theta}^* - \hat{\theta}^*_e\right\| \tag{23}$$

where $\hat{\theta}^*$ is the true optimal, and $\hat{\theta}^*_e$ is the approximate local optimal when we utilize the implicit function theorem to approximate $\Delta\theta$ in Eq. 13.

so that, we have

$$\|\Delta\theta_e\| = \left\|\hat{\theta}^* - \hat{\theta}^*_e\right\| \leq \frac{1}{\lambda} * \left\|\nabla\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \nabla\mathcal{L}(\hat{\theta}^*_e, \hat{\alpha})\right\| \tag{24}$$

Consider the local optimal that $\nabla\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) = 0$, we have

$$\|\Delta\theta_e\| \leq \frac{1}{\lambda} * \left\|\nabla\mathcal{L}(\hat{\theta}^*_e, \hat{\alpha})\right\| \tag{25}$$

Now, we need to bound $\left\|\nabla\mathcal{L}(\hat{\theta}^*_e, \hat{\alpha})\right\|$.

$$\left\|\nabla\mathcal{L}(\hat{\theta}^*_e, \hat{\alpha})\right\| = \|\nabla\mathcal{L}(\theta^* + \Delta\theta, \hat{\alpha})\| \tag{26}$$

Since we utilize Eq. (13) that $\Delta\theta = -\frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta}^{-1} * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta}$ to conduct the second-order Taylor expansion, we have

$$-\frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta} * \Delta\theta = \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta} \tag{27}$$

So

$$
\begin{aligned}
&\|\nabla\mathcal{L}(\theta_e^*,\hat{\alpha})\| = \left\|\nabla_\theta\mathcal{L}(\theta^*+\Delta\theta,\hat{\alpha}) - \nabla_\theta\mathcal{L}(\theta^*,\hat{\alpha}) - \nabla_\theta^2\mathcal{L}(\theta^*,\hat{\alpha}) * \Delta\theta\right\| \\
&= \left\|\int_0^1 \left(\nabla_\theta^2\mathcal{L}(\theta^*+t*\Delta\theta,\hat{\alpha}) - \nabla_\theta^2\mathcal{L}(\theta^*,\hat{\alpha})\right) * \Delta\theta \, dt\right\| \\
&= \left\|\Delta\theta * \int_0^1 \left(\nabla_\theta^2\mathcal{L}(\theta^*+t*\Delta\theta,\hat{\alpha}) - \nabla_\theta^2\mathcal{L}(\theta^*,\hat{\alpha})\right) \, dt\right\| \\
&= \|\Delta\theta\| * \left\|\int_0^1 \left(\nabla_\theta^2\mathcal{L}(\theta^*+t*\Delta\theta,\hat{\alpha}) - \nabla_\theta^2\mathcal{L}(\theta^*,\hat{\alpha})\right) \, dt\right\| \\
&\leq \|\Delta\theta\| * \left\|\int_0^1 C_H * t * \Delta\theta \, dt\right\| = \frac{C_H}{2}\|\Delta\theta\|^2 \\
&= \frac{C_H}{2}\left\|\frac{\partial^2\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta\partial\theta}^{-1} * \frac{\partial\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta}\right\|^2 \\
&\leq \frac{C_H}{2*\sigma_{min}^2}\left\|\frac{\partial\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta}\right\|^2 \\
&\leq \frac{C_H * C_l^2}{2*\sigma_{min}^2}
\end{aligned}
\tag{28}
$$

where the second row is calculated since $\int \nabla_\theta\mathcal{L}^2(\theta^*+t*\Delta\theta,\hat{\alpha}) * \Delta\theta \, dt = \nabla_\theta\mathcal{L}(\theta^*+t*\Delta\theta,\hat{\alpha})$. The fifth row is calculated as we assume that $\mathcal{L}$ is convex and twice-differentiable with C_H. And we assume that the $\frac{\partial\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta}$ is bounded by C_l. σ_{min} is the smallest eigenvalue of Hessian matrix $\frac{\partial^2\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta\partial\theta}$.

In addition, when we consider a more tight bound, and consider a Taylor expansion on α, we have

$$
\begin{aligned}
\left\|\frac{\partial\mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta}\right\|^2 &= \left\|\frac{\partial(\mathcal{L}(\theta^*,\alpha) + \Delta\alpha * \frac{\partial\mathcal{L}(\theta^*,\alpha)}{\partial\alpha} + o(\Delta\alpha))}{\partial\theta}\right\|^2 \\
&\leq \left\|\Delta\alpha\frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}\right\|^2 + o(\|\Delta\alpha\|^2) \\
&\leq C_a^2\|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^2),
\end{aligned}
\tag{29}
$$

where the second row is obtained based on the local optimal $\frac{\partial\mathcal{L}(\theta^*,\alpha)}{\partial\theta} = 0$. So, we have

$$
\|\nabla\mathcal{L}(\theta_e^*,\hat{\alpha})\| \leq \frac{C_H * C_a^2}{2*\sigma_{min}^2} * \|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^2), \tag{30}
$$

and

$$
\|\Delta\theta_e\| \leq \frac{C_H * C_a^2}{2*\sigma_{min}^2 * \lambda} * \|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^2). \tag{31}
$$

Based on Eq. (18), (22), and (31), we have

$$
\|\Delta\theta\| \leq \frac{C_L}{\lambda} * \|\Delta\alpha\| + \frac{C_H * C_a^2}{2*\sigma_{min}^2 * \lambda} * \|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^2) \tag{32}
$$

The final error bound can be calculated as:

$$
E \leq \frac{K^3}{6}\max\left|\frac{\partial\mathcal{L}^3}{\partial\theta^3}\right| \tag{33}
$$

where $K = \frac{C_L}{\lambda} * \|\Delta\alpha\| + \frac{C_H * C_a^2}{2*\sigma_{min}^2 * \lambda} * \|\Delta\alpha\|^2 + o(\|\Delta\alpha\|^2)$.

□

Derivation of Eq. (6)**:** In Section 4, when we consider a first-order Taylor expansion on α as we only apply an infinitesimal change of α, and we alsoconsider a second-order Taylor expansion on θ, we have

$$
\begin{aligned}
\Delta\mathcal{L} =& \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \approx \mathcal{L}(\hat{\theta}^*, \alpha) + \Delta\alpha^T \frac{\partial \mathcal{L}(\hat{\theta}^*, \alpha)}{\partial \alpha} - \mathcal{L}(\theta^*, \alpha) \\
&\approx \Delta\theta^T * \frac{\partial \mathcal{L}(\theta^*, \alpha)}{\partial \theta} + 1/2 \Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} \Delta\theta + \Delta\alpha^T * \frac{\partial \mathcal{L}(\hat{\theta}^*, \alpha)}{\partial \alpha} \\
&= 1/2 \Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} \Delta\theta + \Delta\alpha^T * \frac{\partial \mathcal{L}(\hat{\theta}^*, \alpha)}{\partial \alpha},
\end{aligned} \tag{34}
$$

where the last row is obtained is obtained due to the local optimal $\frac{\partial \mathcal{L}(\theta^*, \alpha)}{\partial \theta} = 0$. $\square$

Theorem 2: Supposed that DARTS obtains the optimized architecture parameter α with supernet weights θ^* after supernet training, and we pose an infinitesimal change on α. Based on implicit function theorem and under the assumption that the third-derivative of the loss function at optimum is zero or sufficiently small [4], the change of validation performance can be estimated as:

$$
\Delta\mathcal{L} = \mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) \approx -1/2 \Delta\alpha^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\alpha\partial\theta} * H^{-1} * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\alpha} * \Delta\alpha, \tag{35}
$$

where $H = \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta}$ is the Hessian matrix.

Proof of Theorem 2: Although Section 3 provides an iterative solution to measure the operation importance based Eq. (5) through individually removing each operation, we need to calculate n times (the number of all candidate operations in the supernet work) to estimate the importance of all candidate operations. More important, since removing one operation poses a considerable change on α ($\alpha_i \to 0$), the approximation error produced by Eq.(5) is non-negligible which may affects the accuracy as stated by Corollary 1. So, a more practical solution to illustrate the importance of each operation is to estimate how the validation performance will change after posing an infinitesimal change on α, a.k.a. operation sensitivity, and we can directly get the operation importance with one calculation.

When we consider that DARTS obtains the optimized architecture parameter α with supernet weights θ^* after supernet training, and we pose an infinitesimal change on α, we can conduct a first-order Taylor expansion on α a second-order Taylor expansion on θ. So that we have get the Eq. (34) as stated in the above. When we further conduct a second Taylor expansion on $\hat{\theta}^*$ for $\frac{\partial \mathcal{L}(\hat{\theta}^*, \alpha)}{\partial \alpha}$, we have

$$
\begin{aligned}
\Delta\mathcal{L} &\approx 1/2 \Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} \Delta\theta + \Delta\alpha^T * \frac{\partial \mathcal{L}(\hat{\theta}^*, \alpha)}{\partial \alpha} \\
&\approx 1/2 \Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} \Delta\theta + \Delta\alpha^T * \left(\frac{\partial(\mathcal{L}(\theta^*, \alpha) + \Delta\theta^T * \frac{\partial \mathcal{L}(\theta^*, \alpha)}{\partial \theta} + \Delta\theta * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} * \Delta\theta^T)}{\partial \alpha}\right) \\
&\approx 1/2 \Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\theta\partial\theta} \Delta\theta + \Delta\alpha^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\alpha\partial\theta} * \Delta\theta,
\end{aligned} \tag{36}
$$

where the last row is obtained when we neglect the third and higher derivatives of the loss function with considering the chain rule that $\frac{\partial \mathcal{L}}{\partial \alpha} = \frac{\partial \mathcal{L}}{\partial \theta} * \frac{\partial(\theta^*(\alpha))}{\partial \alpha}$, and assume α is a local optimal for the continuous magnitude optimization after architecture search. From the above, we can find that, the challenge is still that it is intractable to calculate $\Delta\theta$. Similar as before, when we consider implicit function and second-order Taylor expansion on θ, we have

$$
\frac{\partial \mathcal{L}(\hat{\theta}^*, \hat{\alpha})}{\partial \theta} = 0 \tag{37}
$$

$$
\frac{\partial(\mathcal{L}(\theta^*, \alpha) + \Delta\alpha^T * \frac{\partial \mathcal{L}(\theta^*, \alpha)}{\partial \alpha} + \Delta\theta^T * \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta} + \Delta\theta * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta} * \Delta\theta^T)}{\partial \theta} \approx 0 \tag{38}
$$

When we neglect the third and higher derivatives and the local optimal $\frac{\partial \mathcal{L}(\theta^*, \alpha)}{\partial \theta} = 0$, we have

$$
\Delta\theta^T * \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial\theta\partial\theta} \approx -\Delta\alpha^T * \frac{\partial^2 \mathcal{L}(\theta^*, \alpha)}{\partial\alpha\partial\theta} \tag{39}
$$

Then, similar as before, when we apply a first-order Taylor expansion on $\hat{\alpha}$, with the chain-rule and neglect third and higher derivatives, we have

$$\Delta\theta^T = -\Delta\alpha^T * \frac{\partial^2\mathcal{L}(\theta^*,\alpha^*)}{\partial\alpha\theta} * \frac{\partial^2\mathcal{L}(\theta^*,\alpha^*)}{\partial\theta\partial\theta}^{-1} \tag{40}$$

After applying Eq. (40) on Eq. (36), we have

$$\Delta\mathcal{L} = -1/2\Delta\alpha^T * \frac{\partial^2\mathcal{L}(\theta^*,\alpha^*)}{\partial\alpha\theta} * H^{-1} * \frac{\partial^2\mathcal{L}(\theta^*,\alpha^*)}{\partial\theta\partial\alpha} * \Delta\alpha \tag{41}$$

□

B Practical Implementation

To practically calculate the proposed influential magnitude as Definition 1, we need to estimate the Inverse-Hessian-Vector products (IHVPs) $H^{-1}v$, and $v * \frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}$. In the following, we first described two methods to estimate $H^{-1}v$, and the following DARTS to described the practical calculation on $v * \frac{\partial^2\mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}$.

B.1 Neumann series approximation for Lemma 1

The most common method for the IHVPs is the Neumann series that, for $\|I - A\| < 1$, we have:

$$A^{-1} = \sum_{k=0}^{\infty}(I - A)^k. \tag{42}$$

However, it is impossible to guarantee that $\|I - H\| < 1$. Different from [22], we consider $H^{-1} = \gamma(\gamma H)^{-1}$. When assuming H is bounded with σ_{max}, $\left\|\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right\| < \sigma_{max}$. With $\gamma < \frac{1}{\sigma_{max}}$, we have $\left\|I - \gamma\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right\| < 1$ [29, 36]. When we conduct the Neumann series approximation for $\left[\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right]^{-1}$ in the optimal point, we have:

$$\left[\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right]^{-1} = \gamma(I - I + \gamma\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta})^{-1} = \gamma\sum_{j=0}^{\infty}\left[I - \gamma\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right]^j. \tag{43}$$

$$H^{-1} = \gamma\sum_{k=0}^{\infty}\left[I - \gamma\frac{\partial^2\mathcal{L}}{\partial\theta\partial\theta}\right]^k \approx \gamma\sum_{k=0}^{K}\left[I - \gamma H\right]^k, \tag{44}$$

when we only consider the first K terms and H is positive in θ^* (Please note that, generally, we could only guarantee H is positive when in optimal point). Accordingly, the Lemma 1 is derived.

As suggested by [49], a medium K is large enough to approximate IHVPs, where $K = 2$ obtains similar results as $K \geqslant 3$ while with less computational cost. In this paper, we also observed a similar phenomenon and also set $K = 2$ by default in our Neumann series approximation based IHVPs.

B.2 Sherman-Morrison formula for Lemma 2

Apart from the Neumann series approximation, we could also follow the Sherman-Morrison formula to exactly calculate the inverse of Hessian when we consider the empirical Fisher to replace the Fisher matrix. First, we know that the Hessian equals to fisher information matrix in the optimal point.

$$H = \nabla_\theta\nabla_\theta L = \hat{F} - \frac{1}{n}(y_i - f(x_i))\nabla_\theta\nabla_\theta f \in \mathbb{R}^{p\times p} \tag{45}$$

We name the $r = \frac{1}{n}(y_i - f(x_i))\nabla_\theta\nabla_\theta f$, which is the residual contribution to Hessian H. At the end of training, if we can achieve zero training error, namely, $y_i = f(x_i)$, the residual will converge to zero. We have

$$H = \hat{F} = \nabla_\theta f(\nabla_\theta f)^T. \tag{46}$$

Please note that the dimension of $\nabla_\theta f$ is $M * m$ where m is the number of classes (output dimension) and M is the number of parameters. Practically, we usually consider an empirical Fisher to approximate the Fisher matrix that

$$F = \frac{1}{N}\sum_{n=1}^{N}\nabla_\theta\mathcal{L}(y_n, f(x_n))(\nabla_\theta\mathcal{L}(y_n, f(x_n))^T \tag{47}$$

So, in the remaining text, the F indicates the empirical Fisher for simplicity.

The Sherman-Morrison formula is defined as:

$$(A + \mathrm{u}\mathrm{v}^T)^{-1} = A^{-1} - \frac{A^{-1}\mathrm{u}\mathrm{v}^T A^{-1}}{1 + \mathrm{v}^T A^{-1}\mathrm{u}} \tag{48}$$

When we define that $F_n = \frac{1}{N}\sum_{n=1}^{n} \nabla_\theta \mathcal{L}(y_n, f(x_n))(\nabla_\theta \mathcal{L}(y_n, f(x_n))^T$, we can notice that the empirical Fisher can be represented as F_N, and we have:

$$F_n = F_{n-1} + \frac{1}{N}\nabla_\theta \mathcal{L}_n \, \nabla_\theta \mathcal{L}_n{}^T \tag{49}$$

where $\mathcal{L} = \ell + \eta\mathcal{R}(\theta)$ that ℓ is a cross-entropy loss and $\mathcal{R}$ is the regularization term, and η is the weight-decay, $F_0 = \eta I$. $\mathcal{L}_n$ is the loss for the *n-th* batch dataset. And based on the Sherman-Morrison formula, we have

$$F_{n+1}^{-1} = F_n^{-1} - \frac{F_n^{-1}\nabla_\theta \mathcal{L}_{n+1}\nabla_\theta \mathcal{L}_{n+1}{}^T F_n^{-1}}{N + \nabla_\theta \mathcal{L}_{n+1}{}^T F_n^{-1}\nabla_\theta \mathcal{L}_{n+1}} \tag{50}$$

Please notice that, whether using the Neumann series or the Sherman-Morrison formula, we always need to calculate the Inverse-Hessian Vector Products (IHVPs) rather than only the Inverse-Hessian. It is easy to implement the IHVPs based on the Neumann series through the Hessian-vector products based on Lemma 1. In the following, we discuss how to implement the IHVPs based on Sherman-Morrison formula. When we assume that we have $n+1$ batches datasets, we have

$$H^{-1}v = F_N^{-1}v = F_n^{-1}v - \frac{F_{N-1}^{-1}\nabla_\theta \mathcal{L}_N \nabla_\theta \mathcal{L}_N{}^T F_{N-1}^{-1}}{N + \nabla_\theta \mathcal{L}_N{}^T F_n^{-1}\nabla_\theta \mathcal{L}_N} v \tag{51}$$

We can find that, to solve the previous equation, we need to first recurrently calculate $\mathrm{r}_N = F_{N-1}^{-1}\nabla_\theta \mathcal{L}_N$.

$$F_n^{-1}v = F_{n-1}^{-1}v - \mathrm{r}_n \frac{\mathrm{r}_n{}^T v}{N + \nabla_\theta l_n{}^T \mathrm{r}_n} = \lambda^{-1}v - \sum_{j=1}^{n} \mathrm{r}_j \frac{\mathrm{r}_j{}^T v}{N + \nabla_\theta \mathcal{L}_j{}^T \mathrm{r}_j} \tag{52}$$

We first need to calculate the $\mathrm{r}_n = F_{n-1}^{-1}\nabla_\theta \mathcal{L}_n$.

$$\mathrm{r}_n = F_{n-1}^{-1}\nabla_\theta \mathcal{L}_n = \lambda^{-1}\nabla_\theta \mathcal{L}_n - \sum_{j=1}^{n-1} \mathrm{r}_j \frac{\mathrm{r}_j{}^T \nabla_\theta \mathcal{L}_n}{N + \nabla_\theta \mathcal{L}_j{}^T \mathrm{r}_j}. \tag{53}$$

So, we can first iterative calculate r_n, and then calculate $F_n^{-1}v$.

B.3 Practical Calculation on Matrix-Vector product

Apart from the IHVPs, we also need to calculate $v_a * \frac{\partial^2 \mathcal{L}}{\partial\alpha\partial\theta}$ and $v_m * \frac{\partial^2 \mathcal{L}}{\partial\theta\partial\alpha}$, which are called as matrix-vector product and computational-expensive. Following DARTS [27], we also utilize the finite difference approximation for the practical calculation.

Let ξ be a small scalar, and consider the function $\frac{\partial \mathcal{L}(\theta, \alpha + \xi * v_a)}{\partial\theta}$ with Taylor expansion on α, we have

$$\begin{aligned} \frac{\partial \mathcal{L}(\theta, \alpha + \xi * v_a)}{\partial\theta} &= \frac{\partial \mathcal{L}(\theta, \alpha)}{\partial\theta} + \xi * v_a \frac{\partial^2 \mathcal{L}(\theta, \alpha)}{\partial\alpha\partial\theta} + ..., \\ \frac{\partial \mathcal{L}(\theta, \alpha - \xi * v_a)}{\partial\theta} &= \frac{\partial \mathcal{L}(\theta, \alpha)}{\partial\theta} - \xi * v_a \frac{\partial^2 \mathcal{L}(\theta, \alpha)}{\partial\alpha\partial\theta} + ..., \end{aligned} \tag{54}$$

when we only consider the first two terms, we have

$$v_a * \frac{\partial^2 \mathcal{L}}{\partial\alpha\partial\theta} = \frac{\frac{\partial \mathcal{L}(\theta, \alpha + \xi * v_a)}{\partial\theta} - \frac{\partial \mathcal{L}(\theta, \alpha - \xi * v_a)}{\partial\theta}}{2\xi}. \tag{55}$$

Similarly, the $v_m * \frac{\partial^2 \mathcal{L}}{\partial\theta\partial\alpha}$ can be approximated as:

$$v_m * \frac{\partial^2 \mathcal{L}}{\partial\theta\partial\alpha} = \frac{\frac{\partial \mathcal{L}(\theta + \xi * v_m, \alpha)}{\partial\alpha} - \frac{\partial \mathcal{L}(\theta - \xi * v_m, \alpha)}{\partial\alpha}}{2\xi}. \tag{56}$$

Table 4: Comparison results with NAS baselines on NAS-Bench-201.

Method	CIFAR-10		CIFAR-100		ImageNet-16-120	
	Valid(%)	Test(%)	Valid(%)	Test(%)	Valid(%)	Test(%)
Random baseline	83.20±13.28	86.61±13.46	60.70±12.55	60.83±12.58	33.34±9.39	33.13±9.66
RandomNAS [26]	80.42±3.58	84.07±3.61	52.12±5.55	52.31±5.77	27.22±3.24	26.28±3.09
ENAS [33]	37.51±3.19	53.89±0.58	13.37±2.35	13.96±2.33	15.06±1.95	14.84±2.10
GDAS [10]	89.88±0.33	93.40±0.49	70.95±0.78	70.33±0.87	41.28±0.46	41.47±0.21
SETN [11]	84.04±0.28	87.64±0.00	58.86±0.06	59.05±0.24	33.06±0.02	32.52±0.21
SNAS [42]	90.10±1.04	92.77±0.84	69.69±2.39	69.35±1.98	42.84±1.79	43.16±2.64
PC-DARTS [43]	89.96±0.15	93.41±0.30	67.12±0.39	67.48±0.89	40.83±0.08	41.31±0.22
DARTS (1st) [27]	39.77±0.00	54.30±0.00	15.03±0.00	15.61±0.00	16.43±0.00	16.32±0.00
DARTS (2nd) [27]	39.77±0.00	54.30±0.00	15.03±0.00	15.61±0.00	16.43±0.00	16.32±0.00
DARTS-PT [40]	87.34±0.43	89.63±0.19	62.48±2.89	62.35±2.14	36.35±2.76	36.51±2.13
DARTS-IF	90.13±0.54	91.84±0.84	65.47±1.33	67.94±1.23	42.78±3.57	42.50±3.30
DARTS-IM	**90.92±0.34**	**93.61±0.23**	**71.21±0.55**	**71.31±0.40**	**44.70±0.74**	**44.98±0.36**
optimal	91.61	94.37	74.49	73.51	46.77	47.31

DARTS-IM best single run achieves **94.29**%, **72.67**%, and **45.93**% test accuracy on CIFAR-10, CIFAR-100, and ImageNet, respectively.

C Why perturbation-based method fails

The perturbation-based method, e.g. the leave-one-out retraining [22], is an intuitive way to measure the importance of every candidate operations. More detailed, we can individually remove every candidate operation in each edge, and then fine-tune the supernet to the validation performance drop $\mathcal{L}(\hat{\theta}^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha)$. However, DARTS-PT abandons the retraining part that directly measure the validation performance drop. An concern will be raised here as the pruned supernet is not trained to optimal, and it is inaccurate to measure the performance drop by $\mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha)$. More importantly, since the *skip-connection*'s magnitude is generally much higher than other candidate operations, and directly removing *skip-connections* usually greatly deteriorates the supernet performance if we do not fine-tune the supernet. This is also the reason why we empirically find that DARTS-PT still prefers *skip-connection* than other operations. Figure 8 (a) plots the searched architecture in NAS-Bench-201 for DARTS-PT, which contains 3 *skip-connections*. In contrast, the architecture selected by our influential magnitude only contains 1 *skip-connection*. This phenomenon also exists in the experiments on the search space S1 and S2 of [45], where although DARTS-PT can avoid selecting all *skip-connections* in all edges, its searched architectures still contains intensive *skip-connections*.

D Comparison with SOTAs on NAS-Bench-201

We also compared our DARTS-IM with different NAS baselines on the NAS-Bench-201 search space. Table 4 reports the comprehensive validation and test results on all three datasets, including CIFAR-10, CIFAR-100, and ImageNet. As shown, DARTS is not a valid method for differentiable architecture search, which even leads to poorer results than random baseline. Although considering other supernet training methods can also improve the performance, we can also found that, with only replacing the operation selection metric with our influential magnitude, our DARTS-IM can achieve more competitive results. Moreover, our DARTS-IM achieves near-optimal results, a **94.29**%, **72.67**%, and **45.93**% test accuracy on CIFAR-10, CIFAR-100, and ImageNet, respectively, in a single run. Our DARTS-IF, which follows the perturbation-based selection as DARTS-PT while with an additional approximation term, also outperforms DARTS-PT that again verifies the effectiveness of the influence function explanation. In addition, due to the requirement of fine-tuning in the perturbation-based paradigm, DARTS-PT cost much more than our DARTS-IM. Compared with the perturbation-based selection paradigm, we found that the magnitude paradigm is more efficient and reliable [1].

E Comparison with SOTAs on DARTS space

DARTS search space is a much more complicated search space than NAS-Bench-201 and NAS-Bench-1shot1, and it is intractable to get the ground truth for all architectures in this space. Generally, most existing works only report the best searched architecture, making the reproducibility of architecture search weak in this space. For fair comparison, we also follow the common setting [27, 43] to conduct the architecture search on the

[1] The example codes and training log files could be found in the supplementary material. The trained supernet used in this experiment can be found in https://github.com/anonymous-submission1991/DARTS-IM.

Table 5: Comparison results with state-of-the-art differentiable NAS approaches.

Method	Test Error (%)			Param	$+\times$	Search
	CIFAR-10	CIFAR-100	ImageNet	(M)	(M)	Cost
SETN [11]	2.69	17.25	25.7 / 8.0	4.6	610	1.8
SNAS [42]	2.85±0.02	20.09	27.3 / 9.2	2.8	474	1.5
PARSEC [6]	2.86±0.06	-	26.3	3.6	620	0.6
MdeNAS [51]	2.55	17.61	25.5 / 7.9	3.6	506	0.16
GDAS [10]	2.93	18.38	26.0 / 8.5	3.4	545	0.21
PC-DARTS [43]	2.57±0.07	17.11	25.1 / 7.8	3.6	586	0.3
DARTS (1st) [27]	2.94	17.76	-	2.9	513	0.15†
DARTS [27]	2.76±0.09	17.54	26.9 / 8.7	3.4	574	0.4†
DARTS-PT [27]	2.61±0.08	17.49	26.1 / 8.2	3.3	536	0.8†
DARTS-IM	**2.50±0.10**	**17.02**	**25.0 / 7.6**	3.8	599	0.4†

"Param" is the model size when applied on CIFAR-10, while "$+\times$" (the number of multiply-add operations) is calculated based on the ImageNet dataset. In our experiments, we adjust the number of initial filters to restrict "$+\times$" to be less than 600M on the ImageNet and CIFAR10. † means the computational time is calculated with the same environment as [40]. The unit in "Search Cost" is GPU day.

Table 6: Analyze different approximation method on NAS-Bench-201.

Method	CIFAR-10		CIFAR-100		ImageNet-16-120	
	Valid(%)	Test(%)	Valid(%)	Test(%)	Valid(%)	Test(%)
Random baseline	83.20±13.28	86.61±13.46	60.70±12.55	60.83±12.58	33.34±9.39	33.13±9.66
DARTS [27]	39.77±0.00	54.30±0.00	15.03±0.00	15.61±0.00	16.43±0.00	16.32±0.00
DARTS-PT [40]	87.34±0.43	89.63±0.19	62.48±2.89	62.35±2.14	36.35±2.76	36.51±2.13
DARTS-IM-I	89.89±0.24	93.11±0.17	69.50±0.60	70.17±0.60	44.45±0.73	44.44±0.32
DARTS-IM-D	89.62±1.70	92.82±1.20	69.17±2.61	69.45±2.72	42.16±3.97	42.07±3.50
DARTS-IM-NS	90.05±0.51	93.35±0.38	69.96±1.13	70.26±1.13	44.43±0.95	44.03±0.75
DARTS-IM-SM	**90.92±0.34**	**93.61±0.23**	**71.21±0.55**	**71.31±0.40**	**44.70±0.74**	**44.98±0.36**

CIFAR-10 dataset in this search space, where the best-found cell is repeatedly stacked to form the full structure for evaluation on CIFAR-10, CIFAR-100, and ImageNet datasets. Please note that, rather than using a default 36 for the number of initial filters, we adjust it to restrict the number of multiply-add operations ("$+\times$") to be less than 600M on the ImageNet and to comparable model size with SOTAs on CIFAR10.

Table 5 presents the comparison results with the SOTA differentiable NAS methods. As shown, our DARTS-IM improves the DARTS baselines by a large margin, in terms of test error on CIFAR-10, CIFAR-100, and ImageNet, respectively, again verifying the effectiveness of the proposed influential magnitude. More interesting, instead of changing the complicated supernet training part, we found that simply replacing the magnitude with the proposed influential magnitude could consistently improve the performance of DARTS with only additional hundred seconds computational time. Different from the DARTS-PT which contains the supernet retraining that greatly increases the search time, our DARTS-IM only poses additional hundred seconds in DARTS space. We could also find that, although the best CIFAR-10 test error by DARTS-PT is comparable with our DARTS-IM, our best searched architecture obtains lower test error in CIFAR-100 and ImageNet, showing better transferability.

F Analyse the inverse Hessian approximation

In this section, we detailed discuss the effects of different inverse Hessian approximation methods on the proposed influence magnitude. Apart from the two devised method, Neumann series and Sherman-Morrison formula, we also consider two commonly-used approaches in approximating the inverse Hessian, identity approximation and diagonal approximation, which consider the Hessian matrix as an identity matrix and diagonal matrix, respectively. We conduct a series of experiments on NAS-Bench-201 to analyze different approximation methods, and all comparison results are reported in Table 6. As shown, even with the identity approximation, our influential magnitude $\mathcal{I}_{\mathcal{M}}$ can achieve satisfying results compared with the random baseline and two heuristic methods. This result indicates that the second-order information $\frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta}$ and $\frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha}$ can provide useful information for operation strength estimation. More interesting, when considering the diagonal Hessian information, we found that our DARTS-IM-D could also achieve excellent performance. The results of

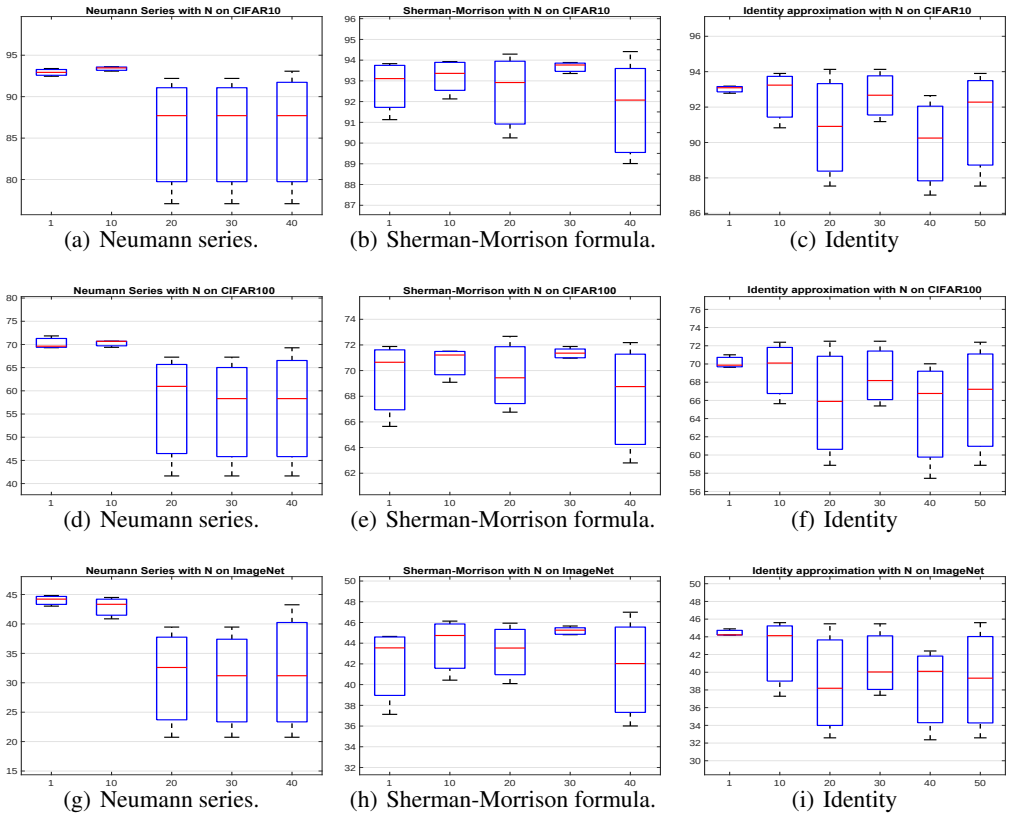

(a) Neumann series. (b) Sherman-Morrison formula. (c) Identity

(d) Neumann series. (e) Sherman-Morrison formula. (f) Identity

(g) Neumann series. (h) Sherman-Morrison formula. (i) Identity

Figure 5: Ablation study on N under two approximation methods, where x-axis is N and y-axis represents test accuracy on CIFAR-10, CIFAR-100, and ImageNet, respectively.

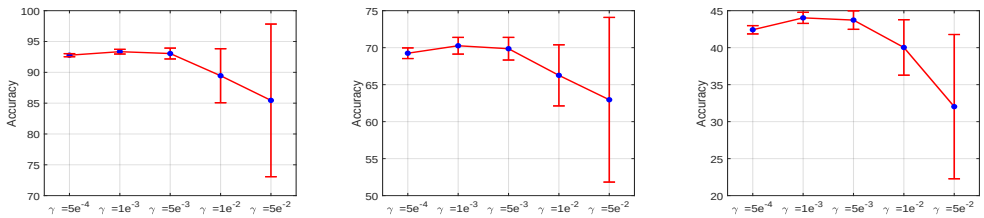

(a) CIFAR-10 (b) CIFAR-100 (c) ImageNet-16-120

Figure 6: Hyperparameter γ analysis of DART-IM-NS on the NAS-Bench-201 benchmark dataset.

identity approximation and diagonal approximation demonstrate the robustness of our influential magnitude, which can still achieve excellent results with different inverse Hessian approximation methods.

However, we also found that an incorrect approximation may lead to poor performance. Similar to Figure 2, we plots the full NAS-Bench-201 results of our influential magnitude under different inverse Hessian matrix approximation methods along with the number of batches N, on the CIFAR10, CIFAR100, and ImageNet, respectively. We compared our DARTS-IM with the DARTS baseline and also the DARTS-PT whose results are directly obtained from [40]. Please note that, DARTS-PT smooth the line extremely with only report the tendency of the architecture selection during the search. We can found that the Neumann series approximation could not obtain stable results, especially when the number of batches increases that the performance of Neumann series even performs more poorly than identity approximation. As discussed in Section 6.1, these results also verify that the stochastic Neumann series estimation for IHVP is somehow erroneous. Generally, we found that the Sherman-Morrison formula is the most practical in estimating the inverse Hessian for our influential magnitude.

Different from Sherman-Morrison formula that only has a hyperparameter N to be tuned and η is fixed in the optimizer, the Neumann series approximation has a key hyperparameter γ [29, 36] to be tuned, which is assumed

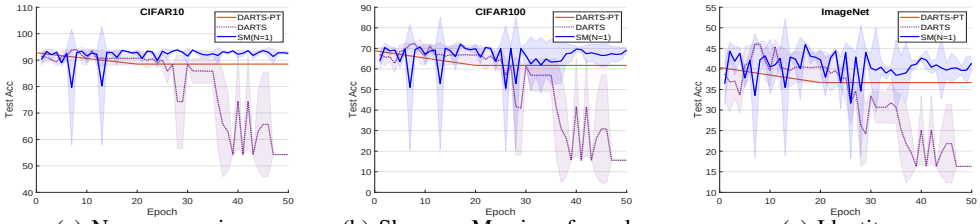

(a) Neumann series. (b) Sherman-Morrison formula. (c) Identity

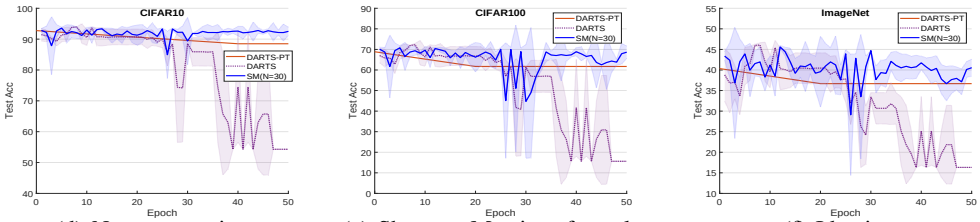

(d) Neumann series. (e) Sherman-Morrison formula. (f) Identity

Figure 7: Track performance of the derived architectures during the search on NAS-Bench-201 with Sherman-Morrison formula under different N for CIFAR-10, CIFAR-100, and ImageNet, respectively.

to be small enough that $\gamma < \frac{1}{\sigma_{max}}$ in Appendix B.1. In this way, we conduct the hyperparameter analysis on γ for our DARTS-IM when employing the Neumann series approximation. Figure 6 compares the performance of DARTS-IM with different γ on the NAS-Bench-201. As shown, the Neumann series approximation is sensitive to the γ, where a large γ significantly deteriorates the performance of DARTS-IM. For example, a $\gamma = 0.05$ even makes our DART-IM similar to the performance of the random baseline. Compared with the Neumann series, we found that the Sherman-Morrison approximation is more robust.

G Visualization of Searched Architectures on Different Spaces

In Figure 8 and 9, we visualize all searched architectures by DARTS baselines, DARTS-PT and also our DARTS-IM on different search space, including NAS-Bench-201 [12], tool search spaces proposed by [45], and DARTS search space [27].

First, we analyze these architectures searched on NAS-Bench-201 [12] and tool search spaces proposed by [40, 45], since which helps us to clearly demonstrate the robustness of different architecture selection methods. As observed by [40, 45], the most notable drawback of DARTS baseline is the robustness and generalization, where DARTS fails to select intensive *skip-connections* with the search progressing. Figure 8 (a-c) plots the searched architectures on NAS-Bench-201, respectively. As shown, the DARTS baseline selects *skip-connection* for all edges, while DARTS-PT can find a valid architecture with only 3 *skip-connections* when considering a perturbation-based selection paradigm. More inspiring, the proposed influential magnitude ($\mathcal{I}_{\mathcal{M}}$) can select more competitive architecture which also only contains one *skip-connection*. Please note that, the trained supernet keeps identical for DARTS, DARTS-PT, and our DARTS-IM, which only adopt different architecture selection methods, e.g., *argmax*, perturbation-based selection, and the proposed influential magnitude.

Similarly, we can also observe this superiority of our $\mathcal{I}_{\mathcal{M}}$ in the tool search spaces proposed by [40, 45], especially in the search space S1 and S2. We can find that, our DARTS-IM clearly outperforms DARTS and DARTS-PT. Although DARTS-PT can partially relieves this instability, it still selects intensive *skip-connections*. In the contrary, our DARTS-IM generally prefers convolutional operations. However, although our DARTS-IM outperforms DARTS baseline in the search space S3, we found that DARTS-PT achieves excellent results in this space. One potential is that this space contains intensive *none* operations which is naturally removed by perturbation-based selection, since removing *none* will not affect the supernet performance. We also conduct the architecture search on the search space S4, while we found that our DARTS-IM also performs poorly as DARTS that select some *noise* operations. One underlying reason is that it is impossible to train the supernet to optimal in this space since the *noise* operation will return a random noise regardless of the input. In addition, DARTS-PT only locally compares the operations' strength in each edge, and a fine-tuning is needed after every step of edge discretization. Differently, our DARTS-IM in this paper consider a global comparison among all operations in a supernet, which may bring some negative operation coupling. To further improve our proposed DARTS-IM, a simple and straight approach is to gradually discretize edges in a supenet with fine-tuning, which we leave

Table 7: Comparison in test error (%) with DARTS baselines on S1-S4 space.

Dataset	Space	DARTS	DARTS-ES	DARTS-PT	DARTS-IM (Avg.)	DARTS-IM (Best)
C10	S1	4.66±0.71	3.05±0.07	3.50	**2.85±0.17**	2.68
	S2	4.42 ± 0.40	3.41 ± 0.14	2.79	**2.51±0.05**	**2.46**
	S3	4.12 ± 0.85	3.71 ± 1.14	**2.49**	3.91±0.07	2.83
	S4	6.95 ± 0.18	4.17 ± 0.21	**2.64**	3.95±0.17	3.72
C100	S1	29.93 ± 0.41	28.90 ± 0.81	24.48	**22.80 ± 0.32**	**22.45**
	S2	28.75 ± 0.92	24.68 ± 1.43	23.16	**21.97 ± 0.61**	**21.27**
	S3	29.01 ± 0.24	26.99 ± 1.79	**22.03**	24.02 ± 0.81	23.12
	S4	24.77 ± 1.51	23.90 ± 2.01	**20.80**	23.58 ± 0.73	22.70
SVHN	S1	9.88 ± 5.50	2.80 ± 0.09	2.62	**2.47 ± 0.12**	2.35
	S2	3.69 ± 0.12	2.68 ± 0.18	2.53	**2.50 ± 0.09**	**2.42**
	S3	4.00 ± 1.01	2.78 ± 0.29	2.42	**2.38 ± 0.11**	**2.28**
	S4	2.90 ± 0.02	2.55 ± 0.15	**2.42**	2.84 ± 0.10	2.73

for our future work. We also show the comparison results in the Table 7 with different DARTS baselines. We directly report the results of DARTS and DARTS-ES from [45], which contain the mean and std. The results of DARTS-PT is also inherited from the original paper [40], which only report the best single run. In Table 7, we report the average and best single run of our DARTS-IM together for better comparison.

Then, we also visualized our searched architectures in the DARTS search space. Interesting, compared with the DARTS baseline, architecture searched by our DARTS-IM contains several *sep_conv* 5×5 operations in the normal cell, where this phenomenon has never been observed by existing popular DARTS based algorithms. Similarly, we also found that our reduction cell contains several *sep_conv* 5×5 operations, which is also in line with another influence based method DARTS-PT and PC-DARTS. More important, different from most existing works that the normal cell generally contains more parameters than the reduction cell, we found that our searched architecture prefers more parameters in the reduction cell. This observation is more consistent with the common experience in the network pruning [24, 25, 39], where we found that most pruning methods will retain more parameters in the reduction cell and prune more redundant parameters in the normal cell.

H Overall Algorithm Framework Description

In this paper, we proposed two different approaches for the operation selection in DARTS, called DARTS-IF (described in Sec.3) and DARTS-IM (described in Sec.4). This two approaches are based on the influence functions to measure the operation importance, where the main different is that DARTS-IF removes a candidate operation in each step to measure the loss change in a perturbation manner, while DARTS-IM calculate the loss sensitivity when applying an infinitesimal change on α. The framework of our DARTS-IF and DART-IM are sketched in Algorithm 1 and 2. Our DARTS-IF shares a similar perturbation paradigm as DARTS-PT, while without any fine-tuning. Rather than using the validation performance drop to indicate the operation importance, DARTS-IF leverages influence functions to predict the loss change based on Eq. (5). In our practical implementation, we consider the Neumann series as described in Lemma 1, or the Sherman-Morrison formula as described in Lemma 2, to approximate the IHPV $\frac{\partial^2 \mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta\partial\theta}^{-1} \frac{\partial \mathcal{L}(\theta^*,\hat{\alpha})}{\partial\theta}$. More important, different from DARTS-PT that gradually discretize edges and finetune the supernet, our DARTS-IF only approximates the effect of independently removing one candidate operation for each edge. In this way, we can get the strength for all candidate operations under every edge in a supernet.

Different from DARTS-IF and DARTS-PT, our DARTS-IM is with a similar paradigm as DARTS, while whose operation importance is calculated based on $\mathcal{I}_{\mathcal{M}}$, rather that the optimized α. In addition, the operation importance in our DARTS-IM can be calculated by one-shot, rather than iteratively going through all candidates operations in a supernet as perturbation paradigm.

In calculating the influence magnitude in our DARTS-IM or the loss change in DARTS-IF, we have three more hyperparameters, γ for Neumann series approximation, η for Sherman-Morrison approximation, and batch size N for both. In our experiments, we set $\gamma = 0.001$ that is same as the learning rate in optimizing θ, and η is the weight decay for regularization, which both can be obtained form the optimizer by default. In the hyperparameter analysis in Figure 2, we found that $N = 10$ is enough for both Neumann series and Sherman-Morrison approximation.

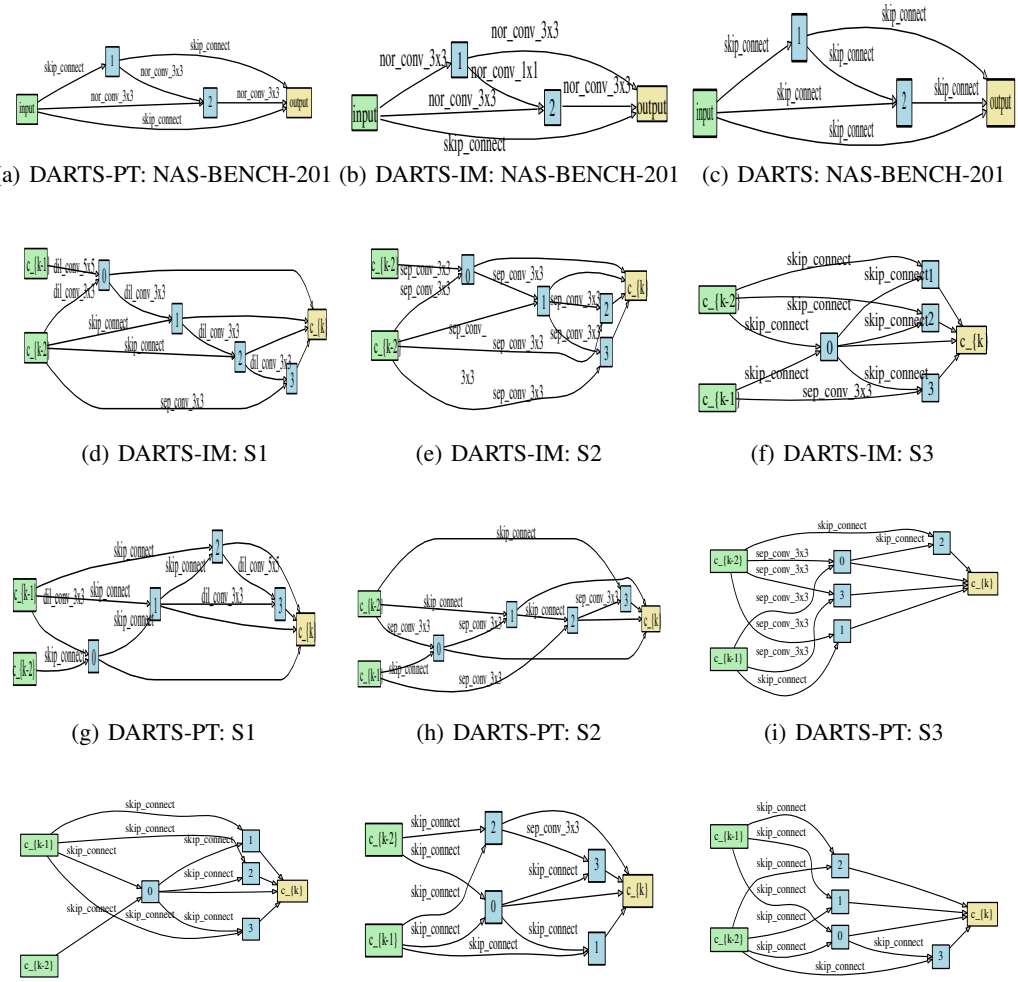

(a) DARTS-PT: NAS-BENCH-201 (b) DARTS-IM: NAS-BENCH-201 (c) DARTS: NAS-BENCH-201

(d) DARTS-IM: S1 (e) DARTS-IM: S2 (f) DARTS-IM: S3

(g) DARTS-PT: S1 (h) DARTS-PT: S2 (i) DARTS-PT: S3

(j) DARTS: S1 (k) DARTS: S2 (l) DARTS: S3

Figure 8: Searched architectures on NAS-Bench-201 and several tool search spaces in [40, 45].

Algorithm 1 N Differentiable Architecture Search with Influence Functions (DARTS-IF)

Input: A pretrained supernet after bi-level training process (θ^*, α), candidate operations for each edge $\mathcal{O}$, and set of edges $\mathcal{E}$ from the supernet.
output: A discrete architecture α^*.
for $e \in \mathcal{E}$ **do**
  for $o \in \mathcal{O}$ **do**
    Remove candidate operation o from edge e;
    Calculate the predictive loss chance $\Delta\mathcal{L}_{o,e}$ based on Eq. (5), that $\Delta\mathcal{L}_{o,e} \approx \mathcal{L}(\theta^*, \hat{\alpha}) - \mathcal{L}(\theta^*, \alpha) - 1/2 \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}^T \frac{\partial^2 \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta \partial \theta}^{-1} \frac{\partial \mathcal{L}(\theta^*, \hat{\alpha})}{\partial \theta}$, as the operation strength;
    Restore o to $\mathcal{O}$;
  end for
end for
Apply *argmax* on the operation strength $\Delta\mathcal{L}$ and derive the discrete architecture α^* accordingly.

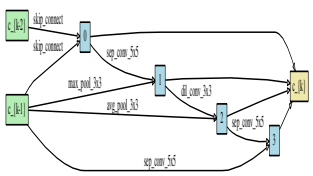

(a) Influential Magnitude ($\mathcal{I}_{\mathcal{M}}$): Normal Cell

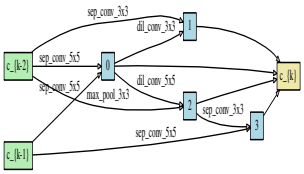

(b) Influential Magnitude ($\mathcal{I}_{\mathcal{M}}$): Reduction Cell

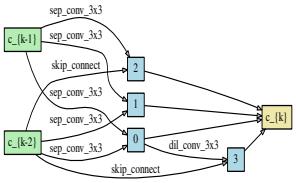

(c) DARTS: Normal Cell

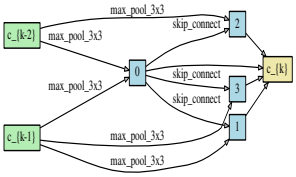

(d) DARTS: Reduction Cell

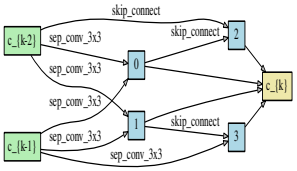

(e) DARTS-PT: Normal Cell

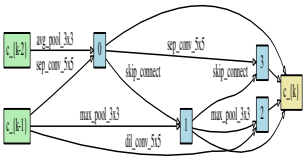

(f) DARTS-PT: Reduction Cell

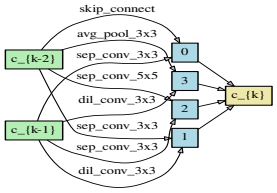

(g) PC-DARTS: Normal Cell

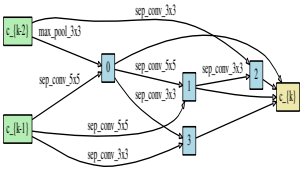

(h) PC-DARTS: Reduction Cell

Figure 9: Searched architectures on DARTS search space.

Algorithm 2 Differentiable Architecture Search with Influence Magnitude (DARTS-IM)

1: **Input**: A pretrained supernet after bi-level training process (θ^*, α), candidate operations for each edge $\mathcal{O}$, and set of edges $\mathcal{E}$ from the supernet.
2: **output:** A discrete architecture α^*.
3: Calculate the influence magnitude $\mathcal{I}_{\mathcal{M}} = -\mathbf{1}^T \frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\alpha\partial\theta} H^{-1} \frac{\partial^2 \mathcal{L}(\theta^*,\alpha)}{\partial\theta\partial\alpha}$ based on Definition1;
4: Apply *argmax* on the influence magnitude $\mathcal{I}_{\mathcal{M}}$ and derive the discrete architecture α^* accordingly.