

Shapley-NAS: Discovering Operation Contribution for Neural Architecture Search

Han Xiao^{1,2}, Ziwei Wang^{1,2}, Zheng Zhu^{1,2}, Jie Zhou^{1,2}, Jiwen Lu^{1,2*}

¹Department of Automation, Tsinghua University, China

²Beijing National Research Center for Information Science and Technology, China

{h-xiao20,wang-zw18}@mails.tsinghua.edu.cn; zhengzhu@ieee.org; {jzhou, lujiwen}@tsinghua.edu.cn

Abstract

In this paper, we propose a Shapley value based method to evaluate operation contribution (Shapley-NAS) for neural architecture search. Differentiable architecture search (DARTS) acquires the optimal architectures by optimizing the architecture parameters with gradient descent, which significantly reduces the search cost. However, the magnitude of architecture parameters updated by gradient descent fails to reveal the actual operation importance to the task performance and therefore harms the effectiveness of obtained architectures. By contrast, we propose to evaluate the direct influence of operations on validation accuracy. To deal with the complex relationships between supernet components, we leverage Shapley value to quantify their marginal contributions by considering all possible combinations. Specifically, we iteratively optimize the supernet weights and update the architecture parameters by evaluating operation contributions via Shapley value, so that the optimal architectures are derived by selecting the operations that contribute significantly to the tasks. Since the exact computation of Shapley value is NP-hard, the Monte-Carlo sampling based algorithm with early truncation is employed for efficient approximation, and the momentum update mechanism is adopted to alleviate fluctuation of the sampling process. Extensive experiments on various datasets and various search spaces show that our Shapley-NAS outperforms the state-of-the-art methods by a considerable margin with light search cost. The code is available at <https://github.com/Euphoria16/Shapley-NAS.git>.

1. Introduction

Neural architecture search (NAS) has attracted great interest in deep learning since it discovers the optimal architecture from a large search space of network components according to task performance and hardware config-

urations. Pioneering works applied reinforcement learning [48], evolutionary algorithms [30, 38], and Bayesian optimization [22] for the architecture search, but the large computational overhead prohibits practical deployment of NAS algorithms. Therefore, it is desirable to design highly efficient search strategies without performance degradation.

To reduce the search cost of architecture search, several efficient search strategies have been presented including one-shot NAS [29], network transformation [3], and architecture optimization [26]. Among these approaches, one-shot NAS preserves the optimal sub-networks from the over-parameterized supernet with weight sharing, which prevents the time-consuming exhaustive training for model evaluation. In particular, DARTS [23] converted the discrete operation selection into continuous mixing weights learning and iteratively optimized the architecture parameters and supernet weights by gradient descent with significantly reduced search cost. However, the magnitude of architecture parameters in DARTS cannot reflect the actual operation importance in general [36, 42, 47]. That is, the operation with the largest parameter magnitude does not necessarily result in the highest validation accuracy, which degrades the performance of derived architectures.

In this paper, we present a Shapley-NAS method to evaluate the operation contribution via the Shapley value of supernet components for neural architecture search. Instead of relying on the magnitude of architecture parameters updated by gradient descent, we consider their practical influences on task performance and propose to directly evaluate their contributions to the validation accuracy. Moreover, we observe that the operations in the supernet are related to each other: combinations of operations might have different joint influences on performance compared with their separate ones. In order to deal with such complex relationships, we leverage Shapley value [31, 32], an important solution to attribute contributions to players in cooperative game theory. Fig. 1 shows the differences between our Shapley-NAS and existing DARTS methods. Shapley value directly measures the contributions of operations according to the vali-

*Corresponding author

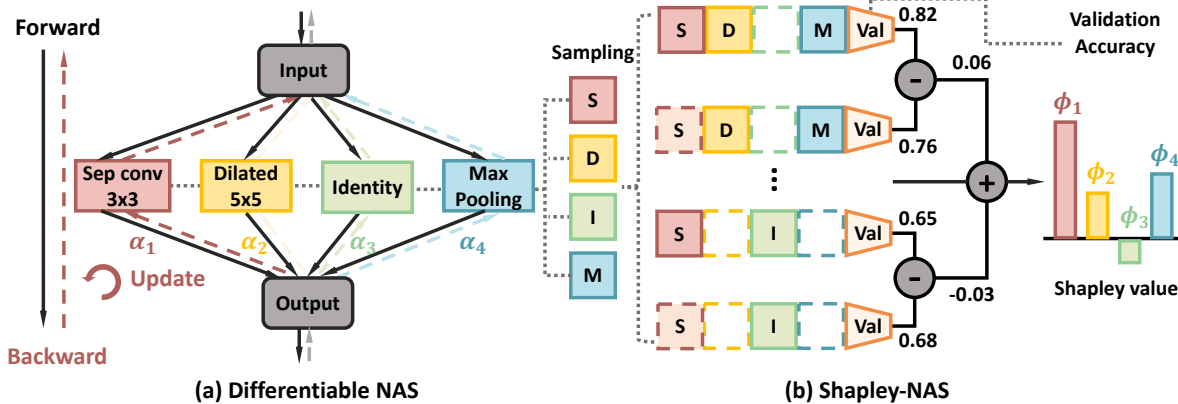


Figure 1. The comparison between DARTS and our Shapley-NAS. (a) DARTS constructs a weight-sharing supernet that consists of all candidate operations. The architecture parameters are optimized by gradient descent, which fails to reflect the importance of operations [36, 42, 47]. (b) The proposed Shapley-NAS method directly evaluates the marginal contribution of operations to the task performance, according to the validation accuracy difference of all possible operation subsets and their counterparts without the given operation.

dation accuracy difference. Meanwhile, it considers all possible combinations and quantifies the average marginal contribution to handle complex relationships between individual elements. Benefiting from these, Shapley value is effective for obtaining operation importance that is highly correlated with task performance. Since computing the exact Shapley value is NP-hard, we employ the Monte-Carlo sampling with early truncation for operation permutation set sampling to approximate it efficiently. Finally, we optimize the supernet weights and update the architecture parameters iteratively, where the momentum update mechanism is adopted to alleviate the fluctuation caused by the sampling process. We empirically demonstrate that the obtained Shapley value has a higher correlation with task performance compared with DARTS. We conducted extensive experiments on different datasets across various search spaces, where our Shapley-NAS outperforms the state-of-the-art architecture search methods. We achieve an error rate of 2.43% on CIFAR-10 [19] on the search space of DARTS and obtain the top-1 accuracy of 23.9% on ImageNet [9] under the mobile setting. Furthermore, our Shapley-NAS acquires the optimal architectures on CIFAR-10 and CIFAR-100 and near-optimal solutions on ImageNet-16-120 of the NAS-Bench-201 benchmark [11].

2. Related Work

Differentiable NAS: Differentiable architecture search (DARTS) was first proposed by Liu *et al.* [23] with the goal of reducing the heavy search cost in NAS. They apply a continuous relaxation to the graphical architecture representation, thus enabling efficient gradient descent to solve the bi-level optimization objective for architecture search. PC-DARTS [40] further proposed to only search the partially-connected operations by leveraging the redundancy in net-

work space to further reduce the memory overhead. Despite the computation efficiency of DARTS, several works have challenged its generalizability [8, 20, 39, 42] and stability [6, 7, 36, 43, 44]. In order to reduce the bias of DARTS for operation selection, SNAS [39] and GDAS [10] introduced stochasticity into the supernet training and adopted the differentiable Gumbel-Softmax trick [17] for gradient estimation. SGAS [20] greedily chose and pruned the candidate operations based on edge importance, selection certainty, and selection stability to alleviate the performance gap between the search and evaluation phase. RobustDARTS [43] found that the stability of DARTS is highly correlated with dominant eigenvalue of the Hessian of validation loss with respect to the architecture parameters. Therefore, they performed early stop regularization according to the largest eigenvalue to avoid poor generation. SmoothDARTS [6] further smoothed the loss landscape via perturbation-based regularization. However, recent studies [36, 42, 47] have demonstrated the magnitude of architecture parameters in the DARTS framework fails to reveal the actual operation importance, which greatly degrades the performance of architectures derived from the search phase.

Shapley value: Shapley value has been well studied in the cooperative game theory as a fair contribution distribution method [31, 32]. Recently, Shapley value has been adopted in explainable machine learning to discover the importance of different elements, which can be divided into three groups: explaining feature importance [1, 24, 25, 28, 33], model component importance [2, 13, 37], and data importance [18, 41]. For the first regard, Ancona *et al.* [1] conducted an axiomatic comparison to show the advantage of the Shapley value over the attribution methods for feature map explanation in deep networks. SHAP [25] presented the additive feature attribution based on the Shapley value of features to acquire higher consistency with hu-

man intuition. For model component importance explanation, ShapNets [37] leveraged the Shapley transform that transforms the input into Shapley representations so that the network prediction can be explained during the forward pass. Neuron Shapley [13] identified the most important filters in neural networks and demonstrated potential applications to improve the accuracy, fairness, and robustness of the model prediction. Moreover, Ghorbani *et al.* [12] quantified the contribution of individual data points which effectively identified the outliers and corrupted data. Since computing the exact Shapley value is NP-hard, Monte-Carlo sampling [12,13], perturbation-based approximation [1], influence function, and many others were presented for efficient estimation of Shapley value. In this paper, we extend the Shapley value to operation importance evaluation in the DARTS framework, so that the optimal architectures are derived by selecting the operations that contribute significantly to the performance.

3. Approach

In this section, we first briefly introduce differentiable architecture search (DARTS), which suffers from degenerate architectures due to the mismatch between the architecture parameters and operation importance. Then we propose to directly evaluate the influence of operations on the task performance and introduce Shapley value to quantify their relative contributions at the presence of complex relationships between different operations. We also present the Monte-Carlo sampling algorithm with early truncation for efficient approximation of Shapley value. Finally, we propose Shapley-based architecture search (Shapley-NAS) which can effectively identify the optimal architectures with the most important operations from the large search space.

3.1. Preliminaries

The differentiable architecture search (DARTS) is one of the most popular solutions to identify effective architectures, as it largely reduces the search cost by relaxing the architecture search to continuous mixture weights learning. Following prior works [22, 30, 49], DARTS searches for the best cell structure and constructs the supernet by repetitions of normal and reduction cells. Each cell is represented by a directed acyclic graph (DAG) with \mathcal{N} nodes and \mathcal{E} edges, where each node $x^{(i)}$ defines a latent representation and each edge (i, j) is associated with an operation $o^{(i,j)}$. The core idea of DARTS is to apply continuous relaxation to the search space to perform the gradient-based search. Concretely, the intermediate node is computed as a softmax mixture of candidate operations:

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

where \mathcal{O} is the set of all candidate operations and $\alpha_o^{(i,j)}$ denotes the mixing weight of operation $o^{(i,j)}$ to construct the supernet. With such relaxation, the architecture search can be performed by jointly optimizing the network weight w and architecture parameters α in a differentiable manner with the following bi-level objective:


$$\min_{\alpha} \mathcal{L}_{val}(w^*, \alpha) \quad \text{s.t.} \quad w^* = \arg \min_w \mathcal{L}_{train}(w, \alpha). \quad (2)$$

During the search stage, the weight-sharing supernet containing all these candidate operations is optimized by gradient descent. At the end of the search stage, the final architecture is derived by selecting the operation with the largest architecture parameter α on every edge across all operation choices, $o^{(i,j)} = \arg \max_{o \in \mathcal{O}} \alpha_o^{(i,j)}$.

3.2. Operation Importance Evaluation

The magnitude-based architecture selection process in DARTS relies on an important assumption that the magnitude of architecture parameters represents the operation importance. In other words, it supposes that operations with low magnitude of α result in weak feature representations and thus have little contribution to the network performance. However, recent studies [36, 42, 47] have shown that the value of architecture parameters does not necessarily reflect the actual operation contribution. In many cases, the operation with the largest α does not result in the highest validation accuracy. Therefore, selecting the best operations based on values of α may lead to significant performance degradation at the evaluation phase.

To solve this problem, we propose to perform the architecture search by identifying operations that contribute the most to the validation accuracy. [36] performs a similar evaluation of operation contribution by removing the target operation from the supernet while keeping all other operations to obtain the performance drop. However, we observe that the operations in the supernet are not independent of each other. To demonstrate the underlying relationships between operations on different edges, we remove only one operation separately on the 4th edge and the 5th edge from the supernet pretrained on the NAS-Bench-201 space, and re-evaluate the supernet accuracy. As shown in Fig. 2a, removing the *skip_connect* operation on the 4th edge and the *conv_3x3* on the 5th edge leads to the most dramatic performance drop. However, we find the impacts of combinations of the two edges differ from the simple accumulation of their separate influence. We remove one operation for both edges simultaneously and enumerate all candidate operation combinations to show the results. As Fig. 2b illustrates, removing *conv_3x3* on the 4th edge and *conv_1x1* on the 5th edge results in the most significant degradation, while removing the combination of *skip_connect* and *conv_3x3* only lead to 3.88% performance drop.



(a) Removing one operation on each edge (b) Removing one operation on both edges at the same time

Figure 2. The performance drop caused by (a) removing the target operation on the 4th and 5th edge separately (b) removing one operation on both edges at the same time and enumerating all combinations.

This observation reveals the complex relationships between different operations on different edges: some operations can collaborate with each other and thus have a significant joint contribution to the supernet performance. To deal with such relationships, we leverage Shapley value [31, 32], an important solution from cooperative game theory, to evaluate the individual contribution. Specifically, the differentiable architecture search process can be uniquely mapped into a cooperative game, providing a practical scheme to quantify the operation contribution based on Shapley value. In a cooperative game, N players associate with each other, and a value function V maps each subset of players $S \subseteq N$ to a real value $V(S)$, which represents the expected payoff the players can obtain by cooperation. In differentiable NAS, the supernet is composed of several layers with identical cell structures, and each cell has $|\mathcal{E}|$ edges each with $|\mathcal{O}|$ operations. Therefore, a set of individual operations, $N = \mathcal{O} \times \mathcal{E} = \{o^{(i,j)}\}_{o \in \mathcal{O}, (i,j) \in \mathcal{E}}$, can be modeled as players in the cooperative game, where all players work together towards the supernet’s performance $V(N)$. Shapley value is utilized to distribute the total performance gains $V(N)$ to each player in N . In our problem, for operation $o^{(i,j)}$, its Shapley value $\phi_o^{(i,j)}$ can be computed as:

$$\phi_o^{(i,j)}(V) = \frac{1}{|N|} \sum_{S \subseteq N \setminus \{o^{(i,j)}\}} \frac{V(S \cup \{o^{(i,j)}\}) - V(S)}{\binom{|N|-1}{|S|}} \quad (3)$$

The Shapley value represents the average marginal contribution of the operation to the network performance, which is obtained by evaluating the performance difference between all operation subsets and their counterparts without the given operation. Here we use the validation accuracy as the value function V to measure the network performance. It has been proved that the formulation of Shapley value in Eq. (3) makes it the only method to quantify individual contributions that uniquely satisfies the following properties [32], which we interpret according to our problem:

Efficiency The performance of the entire supernet is the sum of contributions of individual operations, i.e.

$$\sum_{o^{(i,j)} \in N} \phi_o^{(i,j)} = V(N).$$

Null Player If the operation has no impact on the performance when added to or removed from any subsets of the supernet, then its contribution is zero. That is, if $V(S) = V(S \cup \{o^{(i,j)}\})$ for any operation subset $S \subseteq N \setminus \{o^{(i,j)}\}$, we can derive $\phi_o^{(i,j)} = 0$. For example, the zero operation in the search space of DARTS has no impact on the final performance and thus has zero contribution.

Symmetry If two different operations could be exchanged without affecting the performance, they should be assigned with equal contributions. For any operation subset $S \subseteq N \setminus \{o^{(i,j)}, o^{(k,l)}\}$, $V(S \cup \{o^{(i,j)}\}) = V(S \cup \{o^{(k,l)}\})$, then we have $\phi_o^{(i,j)} = \phi_{o'}^{(k,l)}$.

3.3. Shapley Value Approximation

Although Shapley value can be considered as a desirable attribution metric for quantifying the contribution of operations, directly computing Shapley value from Eq. (3) requires $2^{|\mathcal{O}| \times |\mathcal{E}|}$ network evaluations caused by enumerating all possible subsets. Therefore, the exact computation of Shapley value becomes expensive since $|\mathcal{O}| \times |\mathcal{E}|$ in the common search space is usually large. To efficiently estimate the Shapley value, we present an approximate method based on Monte-Carlo sampling [5]. Specifically, the Shapley value of operation $o^{(i,j)}$ (denoted as o for simplicity) is equivalent to estimating the mean of a random variable, which can be written as:

$$\phi_o(V) = \frac{1}{N!} \sum_{R \in \pi(N)} [V(\text{Pre}_o(R) \cup \{o\}) - V(\text{Pre}_o(R))] \quad (4)$$

where $\pi(N)$ denotes the set of permutations of all elements in N , and $\text{Pre}_o(R)$ is the set of predecessors of o in a given permutation $R \in \pi(N)$. Based on Eq. (4), we can get an unbiased approximation of every operation’s Shapley value by sampling permutations of operation set N . Notably, the Monte-Carlo estimation reduces the exponential calculation complexity to polynomial-time $M \times (|\mathcal{O}| \times |\mathcal{E}|)$, where M is the number of samples. Although this sampling-based estimation requires repetitions of accuracy evaluation on the validation set, it only includes the forward process through the supernet with no need for back-propagation, thus enabling efficient approximation of Shapley value.

Moreover, we find when the number of operations in $\text{Pre}_o(R)$ becomes too small, the task performance degrades dramatically and yields unstable sampling results. Therefore, to reduce the fluctuation of Shapley value estimation, we utilize the early truncation technique during the Monte-Carlo sampling procedure. Specifically, when the masked out operations lead to an extreme performance drop exceeding a pre-defined threshold η , we break off the current sam-

pling. This early truncation technique also reduces nearly half of computation cost, which makes the overall computational overheads comparable with gradient-based architecture parameter optimization in DARTS.

3.4. Shapley-based Architecture Search

We leverage the Shapley value of operations to guide the architecture search to find the best solutions as it reveals the actual operation contribution to performance. Fig. 1 shows the difference between our Shapley-NAS and conventional differential NAS. Instead of updating the architecture parameters by gradient descent, we utilize the Shapley value to represent the relative strength of operations. Specifically, the search objective should be modified as follows:

$$\alpha \propto \phi(\mathcal{L}_{val}(w^*, \alpha)) \quad \text{s.t.} \quad w^* = \arg \min_w \mathcal{L}_{train}(w, \alpha). \quad (5)$$

Since solving the above problem exactly is impractical, we optimize this objective via an approximate way. We update α according to the Shapley value estimated by the algorithm presented in Sec. 3.3:

$$\alpha_t = \alpha_{t-1} + \epsilon \cdot \frac{s_t}{\|s_t\|_2} \quad (6)$$

where α_t means the architecture parameter at the t -th step during the optimization, s_t represents the accumulated Shapley value in the t_{th} step, $\|\cdot\|_2$ is the L_2 norm and ϵ is the step size. We iteratively optimize w_t by descending $\nabla \mathcal{L}_{train}(w_{t-1}, \alpha_{t-1})$ and update architecture parameters α until convergence. To reduce undesired fluctuation in updating caused by random sampling, we introduce the momentum into the iteration to stabilize the optimization:

$$s_t = \mu \cdot s_{t-1} + (1 - \mu) \cdot \frac{\phi(\text{Acc}_{val}(w_{t-1}, \alpha_{t-1}))}{\|\phi(\text{Acc}_{val}(w_{t-1}, \alpha_{t-1}))\|_2} \quad (7)$$

where μ is the momentum coefficient that balances the accumulated Shapley value and the current sampling result, Acc_{val} is the validation accuracy used as value function, and w_{t-1} is the supernet weights at $(t - 1)$ step. After the search stage, we derive the final architecture by selecting the operation with the largest contribution on each edge.

4. Experiments

In this part, we conducted extensive experiments to evaluate our method on the DARTS search space with CIFAR-10 [19] and ImageNet [9] for image classification, as well as on a widely used NAS benchmark dataset, NAS-Bench-201 [11]. We first introduce the datasets and implementation details of our Shapley-NAS. In the following ablation study, we analyzed the effectiveness of the proposed Shapley value evaluation, as well as the influence of hyperparameters on task performance and search cost. We compare our

Shapley-NAS with the state-of-the-art methods with respect to the accuracy, model complexity, and search cost. Finally, we empirically demonstrated the effectiveness of Monte-Carlo sampling estimation, as well as the high correlation between obtained Shapley value and task performance.

4.1. Datasets and Implementation Details

CIFAR-10: For the CNN search space on CIFAR-10, we employed the same operation space \mathcal{O} as DARTS and set the initial channel number as 16. We utilized the partial connection strategy in PC-DARTS [40] to reduce memory overhead and increase batch size. We trained the supernet for 50 epochs (the first 15 epochs for warm-up) with a batch size of 256. The training set of CIFAR-10 was divided into two parts with equal size, one for optimizing network weights and the other for evaluating Shapley value. We set the number of samples M to be 10 in the Monte-Carlo sampling and the early truncation threshold η to be 0.5. The momentum coefficient μ and step size ϵ were assigned to 0.8 and 0.1 respectively. At the evaluation phase, We simply followed the DARTS experimental settings for fair comparison and retrained the network from scratch for 600 epochs.

ImageNet: ImageNet contains about 1.2 million training and 50K validation images from 1000 categories, which is much more challenging than CIFAR-10. We randomly sampled 10% and 2.5% images from the entire 1.3M training set of ImageNet for training network weights and estimating Shapley value respectively. The supernet was trained for 50 epochs with batch size 1024 and the architecture parameters remained frozen in the first 25 epochs. The other hyper-parameters were the same with CIFAR-10. At the evaluation stage, we trained the network from scratch for 250 epochs by an SGD optimizer with a linearly decayed learning rate initialized as 0.5, a momentum of 0.9, and a weight decay of 3×10^{-5} .

NAS-Bench-201: NAS-Bench-201 is a popular benchmark to analyze NAS algorithms, as it provides performance of all candidate architectures which can be directly obtained by querying. In the search space of NAS-Bench-201, the operation set \mathcal{O} has 5 elements and each cell contains 4 nodes, leading to a total search space of 15,625 architectures. NAS-Bench-201 supports three datasets, CIFAR-10, CIFAR-100, and ImageNet-16-120, and more details about the datasets can be found in their paper [11]. Specifically, we acquired the task-specific performance by directly searching on the evaluation dataset, and obtained the mean and standard deviation for the best architecture from 4 independent runs with different random seeds.

4.2. Ablation Study

Effectiveness of Shapley value evaluation: To verify the effectiveness of Shapley-NAS, we conducted experiments on 4 simplified search spaces S1-S4 proposed by


Table 1. The test error(%) of different search algorithms on S1-S4. DARTS+Shapley denotes the combination of DARTS and Shapley value evaluation, and * means freezing α during the search.

Dataset	Method	S1	S2	S3	S4
C10	DARTS	3.84	4.85	3.34	7.20
	DARTS+Shapley	3.11	2.92	2.58	3.45
	DARTS+Shapley*	2.95	2.84	2.67	2.94
	Shapley-NAS	2.82	2.55	2.42	2.63
C100	DARTS	29.46	26.05	28.90	22.85
	DARTS+Shapley	28.21	24.51	23.67	22.78
	DARTS+Shapley*	25.24	24.66	22.39	22.15
	Shapley-NAS	23.60	22.77	21.92	21.53
SVHN	DARTS	4.58	3.53	3.41	3.05
	DARTS+Shapley	2.59	2.72	2.83	2.65
	DARTS+Shapley*	2.88	2.64	2.49	2.58
	Shapley-NAS	2.36	2.43	2.34	2.41

[43] on CIFAR-10, CIFAR100, and SVHN. We first combined the proposed Shapley value evaluation method with DARTS (denoted as DARTS+Shapley in Tab. 1, by only applying Shapley value evaluation at the final discretization step, *i.e.* selecting operations based on their Shapley values instead of α . Moreover, we also tested the performance under the same setting but keeping α frozen, denoted as DARTS+Shapley*. As shown in Tab. 1, DARTS achieves competitive results with the proposed Shapley evaluation method, even when α is not optimized in the training. Notably, our Shapley-NAS still outperforms DARTS+Shapley and DARTS+Shapley*, since taking Shapley value into the supernet optimization can further alleviate the problem caused by gradient-based NAS methods.

Influence of samples times M and early truncation threshold η : We also explored the influence of sampling times M and early truncation threshold η in the Monte-Carlo sampling algorithm. The values of sampling times M and early truncation threshold η are significant for accurate Shapley value estimation, which also affect the overall search cost. Fig. 3 shows the test error (%) and search cost (GPU days) on CIFAR-10 with various M and η . Reducing the number of samples results in lower search cost while degrading the performance since the sampling is not enough to make an accurate estimation. However, the estimation accuracy with samples larger than 10 is not sensitive to the number of samples, and we choose $M = 10$ for search efficiency. Meanwhile, medium η also achieves the best accuracy-complexity trade-off as it mitigates the fluctuation of sampling and reduces the search cost.

Impact of momentum coefficient μ and step size ϵ : To investigate the influence of momentum coefficient μ and step size ϵ on test accuracy, we implemented the architecture parameter assignment with different μ and ϵ . The test error range and model parameter cost are demonstrated in Tab. 2, where medium ϵ outperforms other values. Small step sizes fail to achieve the optimal distribution when reaching the maximum update iterations, and large step sizes make the supernet optimization hard to converge.



(a) Varying number of samples M (b) Varying early truncation threshold η

Figure 3. The test error (%) and search cost (GPU days) of the proposed method on CIFAR-10 with (a) different number of samples and (b) various thresholds of early truncation in the Monte-Carlo sampling for Shapley value estimation.

With the increase of μ , the training stabilization becomes enforced, where μ with 0.8 achieves the best accuracy.

4.3. Comparison with the State-of-the-art NAS Methods

Tab. 3 shows the performance of Shapley-NAS on CIFAR-10 compared with the state-of-the-art NAS methods. Our Shapley-NAS achieves an average test error of 2.47% while only using 0.3 GPU days, significantly surpassing the DARTS baseline in both search cost and accuracy. The test error of the best single run in our experiments is 2.43%, ranking top amongst popular NAS methods. Although ProxylessNAS [4] achieves a lower test error of 2.08%, it performs architecture search on a different space with heavy search cost. The low variance of the experimental results also demonstrates the stability of the proposed search method.

The comparison results on ImageNet with other methods are demonstrated in Tab. 4. We follow the mobile setting in [23] for ImageNet, where the number of multiply-add operations (“ $\times+$ ”) is restricted to be less than 600M. We trained the best-found architecture on CIFAR-10 to evaluate its transferability to ImageNet and obtained a competitive result with 24.3%/7.3% top-1/5 test error, verifying the generalization ability of our Shapley-NAS. We also evaluated the optimal architecture directly searched on ImageNet and obtain a top-1/5 test error of 23.9%/7.2%, which outperforms all other NAS methods with light search cost. Notably, despite the outstanding performance of DrNAS, it has a number of multiply-add operations much over 600M. By contrast, our Shapley-NAS never violates the mobile setting while achieving competitive results.

For NAS-Bench-201, our Shapley-NAS achieves outstanding performance with 94.37%, 73.51%, and 46.85% test accuracy on CIFAR-10, CIFAR-100, and ImageNet-16-120 respectively, as shown in Tab. 5. Notably, we obtain the global optimal architectures on CIFAR-10 and CIFAR-100, which indicates that the proposed method can identify im-

Table 2. The test error (%) and parameter storage cost (M) of the final architectures w.r.t. different values of momentum coefficient μ and different assignments of step size ϵ .

step size ϵ	$\mu = 0.2$		$\mu = 0.5$		$\mu = 0.8$		$\mu = 0.9$	
	Test Error(%)	Params(M)	Test Error(%)	Params(M)	Test Error(%)	Params(M)	Test Error(%)	Params(M)
0.01	2.89 \pm 0.21	4.0	2.87 \pm 0.16	3.7	2.67 \pm 0.06	3.5	2.74 \pm 0.11	3.8
0.05	2.85 \pm 0.18	3.6	2.79 \pm 0.12	3.4	2.55 \pm 0.07	3.2	2.68 \pm 0.07	3.5
0.1	2.82 \pm 0.11	3.7	2.66 \pm 0.10	3.3	2.47 \pm 0.04	3.4	2.61 \pm 0.06	4.1
0.5	2.92 \pm 0.19	3.5	2.84 \pm 0.13	4.2	2.71 \pm 0.12	3.8	2.83 \pm 0.15	3.9

Table 3. Comparison with state-of-the-art image classifiers on CIFAR-10. Means and standard deviations of our Shapley-NAS are obtained by repeated experiments with 4 random seeds.

Architecture	Test Error (%)	Params (M)	Search Cost (GPU days)
DenseNet-BC [16]	3.46	25.6	-
NASNet-A [49]	2.65	3.3	2000
AmoebaNet-A [30]	3.34 \pm 0.06	3.2	3150
AmoebaNet-B [30]	2.55 \pm 0.05	2.8	3150
PNAS [22]	3.41 \pm 0.09	3.2	225
ENAS [29]	2.89	4.6	0.5
NAONet [26]	3.53	3.1	0.4
RandomNAS [21]	2.85 \pm 0.08	4.3	2.7
DARTS (1st order) [23]	3.00 \pm 0.14	3.3	0.4
DARTS (2nd order) [23]	2.76 \pm 0.09	3.3	1.0
SNAS(moderate) [39]	2.85 \pm 0.02	2.8	1.5
GDAS [10]	2.93	3.4	0.3
BayesNAS [46]	2.81 \pm 0.04	3.4	0.2
ProxylessNAS [4]	2.08	5.7	4.0
P-DARTS [8]	2.50	3.4	0.3
PC-DARTS [40]	2.57 \pm 0.07	3.6	0.1
SGAS (Cri 1. avg) [20]	2.66 \pm 0.24	3.7	0.25
SDARTS-RS [6]	2.61 \pm 0.02	3.4	0.4
DrNAS [7]	2.54 \pm 0.03	4.0	0.4
DARTS+PT [36]	2.61 \pm 0.08	3.0	0.8
Shapley-NAS(avg.)	2.47 \pm 0.04	3.4	0.3
Shapley-NAS(best)	2.43	3.6	0.3

Table 4. Comparison with state-of-the-art image classifiers on ImageNet under the mobile setting [23]. † indicates the results obtained by searching on ImageNet, otherwise on CIFAR-10.

Architecture	Test Error (%)	Params (M)	$\times +$ (M)	Search Cost (GPU days)
Inception-v1 [34]	30.1	6.6	1448	-
MobileNet [15]	29.4	4.2	569	-
ShuffleNet 2 \times (v1) [45]	26.4	~ 5	524	-
ShuffleNet 2 \times (v2) [27]	25.1	~ 5	591	-
NASNet-A [49]	26.0	5.3	564	2000
AmoebaNet-C [30]	24.3	6.4	570	3150
PNAS [22]	25.8	5.1	588	225
MnasNet-92 [35]	25.2	4.4	388	-
DARTS (2nd) [23]	26.7	4.7	574	1.0
SNAS (mild) [39]	27.3	4.3	522	1.5
GDAS [10]	26.0	5.3	545	0.3
BayesNAS [46]	26.5	3.9	-	0.2
ProxylessNAS (GPU) [4]†	24.9	7.1	465	8.3
P-DARTS [8]	24.4	4.9	557	0.3
PC-DARTS [40]	25.1	5.3	586	0.1
PC-DARTS [40]†	24.2	5.3	582	3.8
SGAS (Cri 1. best) [20]	24.2	5.3	585	0.25
SDARTS-ADV [6]	25.6	6.1	-	0.4
DrNAS [7]†	24.2	5.2	644	3.9
DARTS+PT [36]†	25.5	4.7	538	3.4
Shapley-NAS	24.3	5.1	566	0.3
Shapley-NAS†	23.9	5.4	582	4.2

portant operations and derive the best architecture from the large search space. On the ImageNet-16-120 dataset, we




Figure 4. The evolution of architecture parameters α by estimating Shapley value based on Monte-Carlo Sampling. (a) The curves of α on the first edge of the normal cell. (b) The curves of α on the first edge of the reduction cell.

also acquire a near-optimal solution, which outperforms the state-of-the-art algorithms, again verifying the effectiveness of our Shapley-NAS.

5. Performance Analysis

Shapley value estimation by Monte-Carlo Sampling:

To verify the effectiveness of Monte-Carlo Sampling for Shapley value approximation, we plot the architecture parameters evolution on the first edge of normal and reduction cells in Fig. 4. Note that the curves of the first 15 epochs for warm-up are not presented. As Fig. 4a shows, although the *max_pool_3x3* operation is larger than all other operations at the start, the operation *sep_conv_5x5* finally becomes the strongest operation since it has the most contribution to the supernet along with the training process. While in Fig. 4b, the operation *sep_conv_3x3* becomes dominant after several epochs, while other operations gradually converge to be very weak. The supernet gradually converges to the final derived architecture using the proposed estimation. Moreover, the architecture parameters are differentiated to make the arg max selection more reliable.

Correlation between Shapley value and task performance:

We investigate the correlation between Shapley value of operations and real task performance on NAS-Bench-201. After the search phase, we sample 200 discrete architectures from the search space and compute their corresponding operation strength by averaging the magnitude of architecture parameters. Then we plot the test accuracy obtained by directly querying along with the computed operation strength of DARTS and our Shapley-NAS. We use the Kendall Tau coefficient to measure the correlation, and the results on CIFAR-10, CIFAR-100, and

Table 5. Comparison results with state-of-the-art NAS methods on NAS-Bench-201. † denotes the results are obtained by searching on CIFAR-10, otherwise by directly searching on the evaluation dataset.

Method	CIFAR-10		CIFAR-100		ImageNet-16-120	
	validation	test	validation	test	validation	test
ResNet [14]	90.83	93.97	70.42	70.86	44.53	43.63
Random (baseline)	90.93 ± 0.36	93.70 ± 0.36	70.60 ± 1.37	70.65 ± 1.38	42.92 ± 2.00	42.96 ± 2.15
RSPS [21]	84.16 ± 1.69	87.66 ± 1.69	45.78 ± 6.33	46.60 ± 6.57	31.09 ± 5.65	30.78 ± 6.12
REINFORCE [49]†	91.09 ± 0.37	93.85 ± 0.37	71.61 ± 1.12	71.71 ± 1.09	45.05 ± 1.02	45.24 ± 1.18
ENAS [29]	39.77 ± 0.00	54.30 ± 0.00	10.23 ± 0.12	10.62 ± 0.27	16.43 ± 0.00	16.32 ± 0.00
DARTS [23]†	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 [23]	39.77 ± 0.00	54.30 ± 0.00	38.57 ± 0.00	38.97 ± 0.00	18.87 ± 0.00	18.41 ± 0.00
SNAS [39]	90.10 ± 1.04	92.77 ± 0.83	69.69 ± 2.39	69.34 ± 1.98	42.84 ± 1.79	43.16 ± 2.64
GDAS [10]	90.01 ± 0.46	93.23 ± 0.23	24.05 ± 8.12	24.20 ± 8.08	40.66 ± 0.00	41.02 ± 0.00
PC-DARTS [40]	89.96 ± 0.15	93.41 ± 0.30	67.12 ± 0.39	67.48 ± 0.89	40.83 ± 0.08	41.31 ± 0.22
iDARTS [44]†	89.96 ± 0.60	93.58 ± 0.32	70.57 ± 0.24	70.83 ± 0.48	40.38 ± 0.59	40.89 ± 0.68
DrNAS [7]	91.55 ± 0.00	94.36 ± 0.00	73.49 ± 0.00	73.51 ± 0.00	46.37 ± 0.00	46.34 ± 0.00
Shapley-NAS	91.61 ± 0.00	94.37 ± 0.00	73.49 ± 0.00	73.51 ± 0.00	46.57 ± 0.08	46.85 ± 0.12
optimal	91.61	94.37	73.49	73.51	46.77	47.31




Figure 5. Correlation between operation strength and test accuracy of 200 sampled architectures using DARTS and our Shapley NAS on NAS-Bench-201. The average operation strength is obtained by the magnitude of corresponding architecture parameters in the supernet, and τ is the Kendall Tau coefficient that measures the correlation.

ImageNet-16-120 are shown in Fig. 5. The Shapley value of operations has a higher correlation with the test accuracy ($\tau = 0.526, 0.474, 0.357$ respectively), while the magnitude of α in DARTS is almost entirely uncorrelated with the final task performance. It indicates the effectiveness of Shapley value to help us discover optimal architectures with superior performance during the evaluation phase.

6. Conclusion and Discussion

In this paper, we have presented Shapley-NAS, a Shapley value based operation contribution evaluation method for neural architecture search. Since the architecture parameters updated by gradient descent in DARTS cannot reveal the actual operation importance in general, we propose to directly evaluate the marginal contribution of operations on accuracy via Shapley value. Specifically, the Shapley value of operations can be efficiently approximated by Monte-Carlo sampling based algorithm with early truncation, thus enabling the optimization of the supernet whose architecture parameters are directly updated with the operation contribution. Shapley-NAS achieves state-of-the-art performance on CIFAR-10, ImageNet, and NAS-Bench-

201 benchmarks, which proves its effectiveness to identify the optimal architectures with the most important operations in neural architecture search.

Limitations: The exact computation of Shapley value is expensive on common search space since it needs exponential times of evaluations, and the resulting heavy search burdens would limit the practical applications for task-specific network deployment. Therefore, to evaluate operation contribution during the architecture search, we utilize the Monte-Carlo sampling method which gives an unbiased approximation of Shapley value. Despite being computationally efficient, it might not be as accurate as the exact computation via enumerating all possible subsets.

Acknowledgements

This work was supported in part by the National Key Research and Development Program of China under Grant 2017YFA0700802, in part by the National Natural Science Foundation of China under Grant 62125603 and Grant U1813218, in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI), and in part by a grant from the Institute for Guo Qiang, Tsinghua University.

References

- [1] Marco Ancona, Cengiz Öztireli, and Markus Gross. Explaining deep neural networks with a polynomial time algorithm for shapley value approximation. In *ICML*, pages 272–281, 2019. [2](#), [3](#)
- [2] Marco Ancona, Cengiz Öztireli, and Markus Gross. Shapley value as principled metric for structured network pruning. *arXiv preprint arXiv:2006.01795*, 2020. [2](#)
- [3] Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In *AAAI*, volume 32, 2018. [1](#)
- [4] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018. [6](#), [7](#)
- [5] Javier Castro, Daniel Gómez, and Juan Tejada. Polynomial calculation of the shapley value based on sampling. *Comput. Oper. Res.*, 36(5):1726–1730, 2009. [4](#)
- [6] Xiangning Chen and Cho-Jui Hsieh. Stabilizing differentiable architecture search via perturbation-based regularization. In *ICML*, pages 1554–1565, 2020. [2](#), [7](#)
- [7] Xiangning Chen, Ruochen Wang, Minhao Cheng, Xiaocheng Tang, and Cho-Jui Hsieh. Drnas: Dirichlet neural architecture search. *arXiv preprint arXiv:2006.10355*, 2020. [2](#), [7](#), [8](#)
- [8] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *ICCV*, pages 1294–1303, 2019. [2](#), [7](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. [2](#), [5](#)
- [10] Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *CVPR*, pages 1761–1770, 2019. [2](#), [7](#), [8](#)
- [11] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. *arXiv preprint arXiv:2001.00326*, 2020. [2](#), [5](#)
- [12] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *ICML*, pages 2242–2251, 2019. [3](#)
- [13] Amirata Ghorbani and James Zou. Neuron shapley: Discovering the responsible neurons. *arXiv preprint arXiv:2002.09815*, 2020. [2](#), [3](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [8](#)
- [15] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. [7](#)
- [16] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017. [7](#)
- [17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. [2](#)
- [18] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. Towards efficient data valuation based on the shapley value. In *AISTATS*, pages 1167–1176, 2019. [2](#)
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [2](#), [5](#)
- [20] Guohao Li, Guocheng Qian, Itzel C Delgadillo, Matthias Muller, Ali Thabet, and Bernard Ghanem. Sgas: Sequential greedy architecture search. In *CVPR*, pages 1620–1630, 2020. [2](#), [7](#)
- [21] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *UAI*, pages 367–377, 2020. [7](#), [8](#)
- [22] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *ECCV*, pages 19–34, 2018. [1](#), [3](#), [7](#)
- [23] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. [1](#), [2](#), [6](#), [7](#), [8](#)
- [24] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex De-Grave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable ai for trees. *Nat. Mach. Intell.*, 2(1):56–67, 2020. [2](#)
- [25] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS*, pages 4768–4777, 2017. [2](#)
- [26] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. *arXiv preprint arXiv:1808.07233*, 2018. [1](#), [7](#)
- [27] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. ShuffleNet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, pages 116–131, 2018. [7](#)
- [28] Masayoshi Mase, Art B Owen, and Benjamin Seiler. Explaining black box decisions by shapley cohort refinement. *arXiv preprint arXiv:1911.00467*, 2019. [2](#)
- [29] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *ICML*, pages 4095–4104, 2018. [1](#), [7](#), [8](#)
- [30] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, volume 33, pages 4780–4789, 2019. [1](#), [3](#), [7](#)
- [31] Alvin E Roth. *The Shapley value: essays in honor of Lloyd S. Shapley*. 1988. [1](#), [2](#), [4](#)
- [32] L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, (28):307–317, 1953. [1](#), [2](#), [4](#)
- [33] Erik Strumbelj and Igor Kononenko. An efficient explanation of individual classifications using game theory. *J. Mach. Learn. Res.*, 11:1–18, 2010. [2](#)
- [34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, pages 1–9, 2015. [7](#)
- [35] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnas-

- net: Platform-aware neural architecture search for mobile. In *CVPR*, pages 2820–2828, 2019. 7
- [36] Ruo Chen Wang, Minhao Cheng, Xiangning Chen, Xiaocheng Tang, and Cho-Jui Hsieh. Rethinking architecture selection in differentiable nas. *arXiv preprint arXiv:2108.04392*, 2021. 1, 2, 3, 7
- [37] Rui Wang, Xiaoqian Wang, and David I Inouye. Shapley explanation networks. *arXiv preprint arXiv:2104.02297*, 2021. 2, 3
- [38] Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, and Song Han. Apq: Joint search for network architecture, pruning and quantization policy. In *CVPR*, pages 2078–2087, 2020. 1
- [39] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018. 2, 7, 8
- [40] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pc-darts: Partial channel connections for memory-efficient architecture search. *arXiv preprint arXiv:1907.05737*, 2019. 2, 5, 7, 8
- [41] Gal Yona, Amirata Ghorbani, and James Zou. Who’s responsible? jointly quantifying the contribution of the learning algorithm and data. In *AIES*, pages 1034–1041, 2021. 2
- [42] Kaicheng Yu, Christian Sciuto, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. *arXiv preprint arXiv:1902.08142*, 2019. 1, 2, 3
- [43] Arber Zela, Thomas Elsken, Tonmoy Saikia, Yassine Marrakchi, Thomas Brox, and Frank Hutter. Understanding and robustifying differentiable architecture search. *arXiv preprint arXiv:1909.09656*, 2019. 2, 6
- [44] Miao Zhang, Steven Su, Shirui Pan, Xiaojun Chang, Ehsan Abbasnejad, and Reza Haffari. idarts: Differentiable architecture search with stochastic implicit gradients. *arXiv preprint arXiv:2106.10784*, 2021. 2, 8
- [45] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, pages 6848–6856, 2018. 7
- [46] Hongpeng Zhou, Minghao Yang, Jun Wang, and Wei Pan. Bayesnas: A bayesian approach for neural architecture search. In *ICML*, pages 7603–7613, 2019. 7
- [47] Yuan Zhou, Xukai Xie, and Sun-Yuan Kung. Exploiting operation importance for differentiable neural architecture search. *TNNLS*, 2021. 1, 2, 3
- [48] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 1
- [49] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *CVPR*, pages 8697–8710, 2018. 3, 7, 8