

DeepMAD: Mathematical Architecture Design for Deep Convolutional Neural Network

Xuan Shen^{1*}, Yaohua Wang^{2*}, Ming Lin^{3†}, Yilun Huang², Hao Tang⁴, Xiuyu Sun^{2‡}, Yanzhi Wang^{1‡}

¹Northeastern University, ²Alibaba Group, ³Amazon, ⁴ETH Zurich

{shen.xu,yanzhi.wang}@northeastern.edu, {xiachen.wyh,lielin.hyl,xiuyu.sxy}@alibaba-inc.com,
minglamz@amazon.com, hao.tang@vision.ee.ethz.ch

Abstract

The rapid advances in Vision Transformer (ViT) refresh the state-of-the-art performances in various vision tasks, overshadowing the conventional CNN-based models. This ignites a few recent striking-back research in the CNN world showing that pure CNN models can achieve as good performance as ViT models when carefully tuned. While encouraging, designing such high-performance CNN models is challenging, requiring non-trivial prior knowledge of network design. To this end, a novel framework termed *Mathematical Architecture Design for Deep CNN (DeepMAD¹)* is proposed to design high-performance CNN models in a principled way. In DeepMAD, a CNN network is modeled as an information processing system whose expressiveness and effectiveness can be analytically formulated by their structural parameters. Then a constrained mathematical programming (MP) problem is proposed to optimize these structural parameters. The MP problem can be easily solved by off-the-shelf MP solvers on CPUs with a small memory footprint. In addition, DeepMAD is a pure mathematical framework: no GPU or training data is required during network design. The superiority of DeepMAD is validated on multiple large-scale computer vision benchmark datasets. Notably on ImageNet-1k, only using conventional convolutional layers, DeepMAD achieves 0.7% and 1.5% higher top-1 accuracy than ConvNeXt and Swin on Tiny level, and 0.8% and 0.9% higher on Small level.

1. Introduction

Convolutional neural networks (CNNs) have been the predominant computer vision models in the past decades [23, 31, 41, 51, 62]. Until recently, the emergence of

*These Authors contributed equally.

†Work done before joining Amazon.

‡Corresponding Author

¹Source codes are available at <https://github.com/alibaba/lightweight-neural-architecture-search>

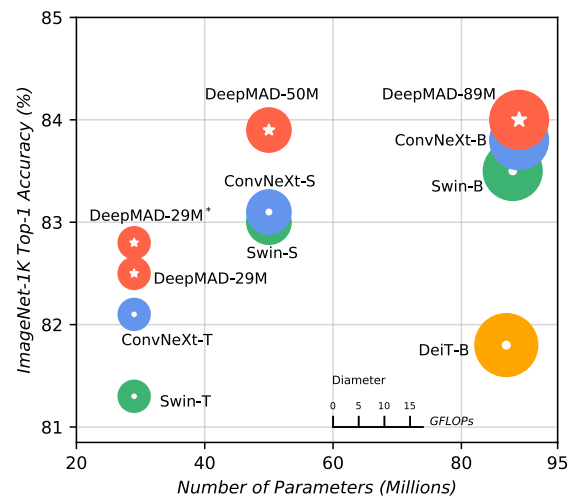


Figure 1. Comparison between DeepMAD models, Swin [40] and ConvNeXt [41] on ImageNet-1k. DeepMAD achieves better performance than Swin and ConvNeXt with the same scales.

Vision Transformers (ViTs) [18, 40, 63] establishes a novel deep learning paradigm surpassing CNN models [40, 63] thanks to the innovation of self-attention [65] mechanism and other dedicated components [3, 17, 28, 29, 54] in ViTs.

Despite the great success of ViT models in the 2020s, CNN models still enjoy many merits. First, CNN models do not require self-attention modules which require quadratic computational complexity in token size [45]. Second, CNN models usually generalize better than ViT models when trained on small datasets [41]. In addition, convolutional operators have been well-optimized and tightly integrated on various hardware platforms in the industry, like IoT [5].

Considering the aforementioned advantages, recent researches try to revive CNN models using novel architecture designs [16, 22, 41, 76]. Most of these works adopt ViT components into CNN models, such as replacing the attention matrix with a convolutional counterpart while keeping the

macrostructure of ViTs. After modifications, these modern CNN backbones are considerably different from the conventional ResNet-like CNN models. Although these efforts abridge the gap between CNNs and ViTs, designing such high-performance CNN models requires dedicated efforts in structure tuning and non-trivial prior knowledge of network design, therefore is time-consuming and difficult to generalize and customize.

In this work, a novel design paradigm named *Mathematical Architecture Design (DeepMAD)* is proposed, which designs high-performance CNN models **in a principled way**. DeepMAD is built upon the recent advances of deep learning theories [8, 48, 50]. To optimize the architecture of CNN models, DeepMAD innovates a constrained mathematical programming (MP) problem whose solution reveals the optimized structural parameters, such as the widths and depths of the network. Particularly, DeepMAD maximizes the differential entropy [26, 32, 58, 59, 67, 77] of the network with constraints from the perspective of *effectiveness* [50]. The *effectiveness* controls the information flow in the network which should be carefully tuned so that the generated networks are well behaved. The dimension of the proposed MP problem in DeepMAD is less than a few dozen. Therefore, it can be solved by off-the-shelf MP solvers nearly instantly on CPU. NO GPU is required and no deep model is created in memory². This makes DeepMAD lightning fast even on CPU-only servers with a small memory footprint. After solving the MP problem, the optimized CNN architecture is derived from the MP solution.

DeepMAD is a mathematical framework to design optimized CNN networks with strong theoretical guarantees and state-of-the-art (SOTA) performance. To demonstrate the power of DeepMAD, we use DeepMAD to optimize CNN architectures only using the conventional convolutional layers [2, 53] as building blocks. DeepMAD achieves comparable or better performance than ViT models of the same model sizes and FLOPs. Notably, DeepMAD achieves 82.8% top-1 accuracy on ImageNet-1k with 4.5G FLOPs and 29M Params, outperforming ConvNeXt-Tiny (82.1%) [41] and Swin-Tiny (81.3%) [40] at the same scale; DeepMAD also achieves 77.7% top-1 accuracy at the same scale as ResNet-18 [21] on ImageNet-1k, which is 8.9% better than He’s original ResNet-18 (70.9%) and is even comparable to He’s ResNet-50 (77.4%). The contributions of this work are summarized as follows:

- A Mathematical Architecture Design paradigm, DeepMAD, is proposed for high-performance CNN architecture design.
- DeepMAD is backed up by modern deep learning theories [8, 48, 50]. It solves a constrained mathematical programming (MP) problem to generate optimized

CNN architectures. The MP problem can be solved on CPUs with a small memory footprint.

- DeepMAD achieves SOTA performances on multiple large-scale vision datasets, proving its superiority. Even only using the conventional convolutional layers, DeepMAD designs high-performance CNN models comparable to or better than ViT models of the same model sizes and FLOPs.
- DeepMAD is transferable across multiple vision tasks, including image classification, object detection, semantic segmentation and action recognition, with consistent performance improvements.

2. Related Works

In this section, we briefly survey the recent works of modernizing CNN networks, especially the works inspired by transformer architectures. Then we discuss related works in information theory and theoretical deep learning.

2.1. Modern Convolutional Neural Networks

Convolutional deep neural networks are popular due to their conceptual simplicity and good performance in computer vision tasks. In most studies, CNNs are usually manually designed [16, 21, 23, 41, 56, 62]. These pre-defined architectures heavily rely on human prior knowledge and are difficult to customize, for example, tailored to some given FLOPs/Params budgets. Recently, some works use AutoML [10, 33, 35, 37, 55, 61, 74] to automatically generate high-performance CNN architectures. Most of these methods are data-dependent and require lots of computational resources. Even if one does not care about the computational cost of AutoML, the patterns generated by AutoML algorithms are difficult to interpret. It is hard to justify why such architectures are preferred and what theoretical insight we can learn from these results. Therefore, it is important to explore the architecture design in a principled way with clear theoretical motivation and human readability.

The Vision Transformer (ViT) is a rapid-trending topic in computer vision [18, 40, 63]. The Swin Transformer [40] improves the computational efficiency of ViTs using a CNN-like stage-wise design. Inspired by Swin Transformer, recent researches combine CNNs and ViTs, leading to more efficient architectures [16, 22, 40, 41, 76]. For example, MetaFormer [76] shows that the attention matrix in ViTs can be replaced by a pooling layer. ConvNext [41] mimics the attention layer using depth-wise convolution and uses the same macro backbone as Swin Transformer [40]. RepLKNet [16] scales up the kernel sizes beyond 31×31 to capture global receptive fields as attention. All these efforts demonstrate that CNN models can achieve as good performance as ViT models when tuned carefully. However, these modern CNNs require non-trivial prior knowledge when designing therefore are difficult to generalize and customize.

²Of course, after solving the MP, training the generated DeepMAD models needs GPU

2.2. Information Theory in Deep Learning

Information theory is a powerful instrument for studying complex systems such as deep neural networks. The Principle of Maximum Entropy [26, 32] is one of the most widely used principles in information theory. Several previous works [8, 50, 52, 59, 77] attempt to establish the connection between the information entropy and the neural network architectures. For example, [8] tries to interpret the learning ability of deep neural networks using subspace entropy reduction. [52] studies the information bottleneck in deep architectures and explores the entropy distribution and information flow in deep neural networks. [77] proposes the principle of maximal coding rate reduction for optimization. [59] designs efficient object detection networks via maximizing multi-scale feature map entropy. The monograph [50] analyzes the mutual information between different neurons in an MLP model. In DeepMAD, the entropy of the model itself is considered instead of the coding rate reduction as in [8]. The *effectiveness* is also proposed to show that only maximizing entropy as in [59] is not enough.

3. Mathematical Architecture Design for MLP

In this section, we study the architecture design for Multiple Layer Perceptron (MLP) using a novel mathematical programming (MP) framework. We then generalize this technique to CNN models in the next section. To derive the MP problem for MLP, we first define the entropy of the MLP which controls its *expressiveness*, followed by a constraint which controls its *effectiveness*. Finally, we maximize the entropy objective function subject to the effectiveness constraint.

3.1. Entropy of MLP models

Suppose that in an L -layer MLP $f(\cdot)$, the i -th layer has w_i input channels and w_{i+1} output channels. The output \mathbf{x}_{i+1} and the input \mathbf{x}_i are connected by $\mathbf{x}_{i+1} = \mathbf{M}_i \mathbf{x}_i$ where $\mathbf{M}_i \in \mathbb{R}^{w_{i+1} \times w_i}$ is trainable weights. Following the entropy analysis in [8], the entropy of the MLP model $f(\cdot)$ is given in Theorem 1.

Theorem 1. *The normalized Gaussian entropy upper bound of the MLP $f(\cdot)$ is*

$$H_f = w_{L+1} \sum_{i=1}^L \log(w_i). \quad (1)$$

The proof is given in Appendix A. The entropy measures the *expressiveness* of the deep network [8, 59]. Following the *Principle of Maximum Entropy* [26, 32], we propose to maximize the entropy of MLP under given computational budgets.

However, simply maximizing entropy defined in Eq. (1) leads to an over-deep network because the entropy grows

exponentially faster in depth than in width according to the Theorem 1. An over-deep network is difficult to train and hinders effective information propagation [50]. This observation inspires us to look for another dimension in deep architecture design. This dimension is termed *effectiveness* presented in the next subsection.

3.2. Effectiveness Defined in MLP

An over-deep network can be considered as a chaos system that hinders effective information propagation. For a chaos system, when the weights of the network are randomly initialized, a small perturbation in low-level layers of the network will lead to an exponentially large perturbation in the high-level output of the network. During the back-propagation, the gradient flow cannot effectively propagate through the whole network. Therefore, the network becomes hard to train when it is too deep.

Inspired by the above observation, in DeepMAD we propose to control the depth of network. Intuitively, a 100-layer network is relatively too deep if its width is only 10 channels per layer or is relatively too shallow if its width is 10000 channels per layer. To capture this relative-depth intuition rigorously, we import the metric termed network *effectiveness* for MLP from the work [50]. Suppose that an MLP has L -layers and each layer has the same width w , the *effectiveness* of this MLP is defined by

$$\rho = L/w. \quad (2)$$

Usually, ρ should be a small constant. When $\rho \rightarrow 0$, the MLP behaves like a single-layer linear model; when $\rho \rightarrow \infty$, the MLP is a chaos system. There is an optimal ρ^* for MLP such that the mutual information between the input and the output are maximized [50].

In DeepMAD, we propose to constrain the *effectiveness* when designing the network. An unaddressed issue is that Eq. (2) assumes the MLP has uniform width but in practice, the width w_i of each layer can be different. To address this issue, we propose to use the average width of MLP in Eq. (2).

Proposition 1. *The average width of an L layer MLP $f(\cdot)$ is defined by*

$$\bar{w} = \left(\prod_{i=1}^L w_i \right)^{1/L} = \exp \left(\frac{1}{L} \sum_{i=1}^L \log w_i \right). \quad (3)$$

Proposition 1 uses geometric average instead of arithmetic average of w_i to define the average width of MLP. This definition is derived from the entropy definition in Eq. (1). Please check Appendix B for details. In addition, geometric average is more reasonable than arithmetic average. Suppose an MLP has a zero width in some layer. Then the information cannot propagate through the network. Therefore, its “equivalent width” should be zero.

In real-world applications, the optimal value of ρ depends on the building blocks. We find that $\rho \in [0.1, 2.0]$ usually gives good results in most vision tasks.

4. Mathematical Architecture Design for CNN

In this section, the definitions of entropy and the *effectiveness* are generalized from MLP to CNN. Then three empirical guidelines are introduced inspired by the best engineering practice. At last, the final mathematical formulation of DeepMAD is presented.

4.1. From MLP to CNN

A CNN operator is essentially a matrix multiplication with a sliding window. Suppose that in the i -th CNN layer, the number of input channels is c_i , the number of output channels is c_{i+1} , the kernel size is k_i , group is g_i . Then this CNN operator is equivalent to a matrix multiplication $W_i \in \mathbb{R}^{c_{i+1} \times c_i k_i^2 / g_i}$. Therefore, the “width” of this CNN layer is projected to $c_i k_i^2 / g_i$ in Eq. (1).

A new dimension in CNN feature maps is the resolution $r_i \times r_i$ at the i -th layer. To capture this, we propose the following definition of entropy for CNN networks.

Proposition 2. For an L -layer CNN network $f(\cdot)$ parameterized by $\{c_i, k_i, g_i, r_i\}_{i=1}^L$, its entropy is defined by

$$H_L \triangleq \log(r_{L+1}^2 c_{L+1}) \sum_{i=1}^L \log(c_i k_i^2 / g_i). \quad (4)$$

In Eq. (4), we use a similar definition of entropy as in Eq. (1). We use $\log(r_{L+1}^2 c_{L+1})$ instead of $(r_{L+1}^2 c_{L+1})$ in Eq. (4). This is because a nature image is highly compressible so the entropy of an image or feature map does not scale up linearly in its volume $O(r_i^2 \times c_i)$. Inspired by [?], taking logarithms can better formulate the ground-truth entropy for natural images.

4.2. Three Empirical Guidelines

We find that the following three heuristic rules are beneficial to architecture design in DeepMAD. These rules are inspired by the best engineering practices.

- **Guideline 1. Weighted Multiple-Scale Entropy** CNN networks usually contain down-sampling layers which split the network into several stages. Each stage captures features at a certain scale. To capture the entropy at different scales, we use a weighted summation to ensemble entropy of the last layer in each stage to obtain the entropy of the network as in [59].
- **Guideline 2. Uniform Stage Depth** We require the depth of each stage to be uniformly distributed as much as possible. We use the variance of depths to measure the uniformity of depth distribution.

- **Guideline 3. Non-Decreasing Number of Channels**

We require that channel number of each stage is non-decreasing along the network depth. This can prevent high-level stages from having small widths. This guideline is also a common practise in a lot of manually designed networks.

4.3. Final DeepMAD Formula

We gather everything together and present the final mathematical programming problem for DeepMAD. Suppose that we aim to design an L -layer CNN model $f(\cdot)$ with M stages. The entropy of the i -th stage is denoted as H_i defined in Eq. (4). Within each stage, all blocks use the same structural parameters (width, kernel size, etc.). The width of each CNN layer is defined by $w_i = c_i k_i^2 / g_i$. The depth of each stage is denoted as L_i for $i = 1, 2, \dots, M$. We propose to optimize $\{w_i, L_i\}$ via the following mathematical programming (MP) problem:

$$\begin{aligned} \max_{w_i, L_i} \quad & \sum_{i=1}^M \alpha_i H_i - \beta Q, \\ \text{s.t.} \quad & L \cdot \left(\prod_{i=1}^L w_i \right)^{-1/L} \leq \rho_0, \\ & \text{FLOPs}[f(\cdot)] \leq \text{budget}, \\ & \text{Params}[f(\cdot)] \leq \text{budget}, \\ & Q \triangleq \exp[\text{Var}(L_1, L_2, \dots, L_M)], \\ & w_1 \leq w_2 \leq \dots \leq w_L. \end{aligned} \quad (5)$$

In the above MP formulation, $\{\alpha_i, \beta, \rho_0\}$ are hyper-parameters. $\{\alpha_i\}$ are the weights of entropies at different scales. For CNN models with 5 down-sampling layers, $\{\alpha_i\} = \{1, 1, 1, 1, 8\}$ is suggested in most vision tasks. Q penalizes the objective function if the network has non-uniform depth distribution across stages. We set $\beta = 10$ in our experiments. ρ_0 controls the *effectiveness* of the network whose value is usually tuned in range $[0.1, 2.0]$. The last two inequalities control the computational budgets. This MP problem can be easily solved by off-the-shelf solvers for constrained non-linear programming [4, 44].

5. Experiments

Experiments are developed at three levels. First, the relationship between the model accuracy and the model *effectiveness* is investigated on CIFAR-100 [30] to verify our effective theory in Section 4.3. Then, DeepMAD is used to design better ResNets and mobile networks. To demonstrate the power of DeepMAD, we design SOTA CNN models using DeepMAD with the conventional convolutional layers. Performances on ImageNet-1K [15] are reported with comparison to popular modern CNN and ViT models.

Finally, the CNN models designed by DeepMAD are transferred to multiple down-streaming tasks, such as MS COCO [36] for object detection, ADE20K [81] for semantic segmentation and UCF101 [57] / Kinetics400 [27] for action recognition. Consistent performance improvements demonstrate the excellent transferability of DeepMAD models.

5.1. Training Settings

Following previous works [71,73], SGD optimizer with momentum 0.9 is adopted to train DeepMAD models. The weight decay is $5e-4$ for CIFAR-100 dataset and $4e-5$ for ImageNet-1k. The initial learning rate is 0.1 with batch size of 256. We use cosine learning rate decay [43] with 5 epochs of warm-up. The number of training epochs is 1,440 for CIFAR-100 and 480 for ImageNet-1k. All experiments use the following data augmentations [47]: mix-up [79], label-smoothing [60], random erasing [80], random crop/resize/flip/lighting, and Auto-Augment [14].

5.2. Building Blocks

To align with ResNet family [21], Section 5.4 uses the same building blocks as ResNet-50. To align with ViT models [18,40,63], DeepMAD uses MobileNet-V2 [23] blocks followed by SE-block [24] as in EfficientNet [62] to design high performance networks.

5.3. Effectiveness on CIFAR-100

The *effectiveness* ρ is an important hyper-parameter in DeepMAD. This experiment demonstrate how ρ affects the architectures in DeepMAD. To this end, 65 models are randomly generated using ResNet blocks, with different depths and widths. All models have the same FLOPs (0.04G) and Params (0.27M) as ResNet-20 [25] for CIFAR-100. The *effectiveness* ρ varies in range [0.1, 1.0].

These randomly generated models are trained on CIFAR-100. The *effectiveness* ρ , top-1 accuracy and network entropy for each model are plotted in Figure 2. We can find that the entropy increases with ρ monotonically. This is because the larger the ρ is, the deeper the network is, and thus the greater the entropy as described in Section 3.1. However, as shown in Figure 2, the model accuracy does not always increase with ρ and entropy. When ρ is small, the model accuracy is proportional to the model entropy; when ρ is too large, such relationship no longer exists. Therefore, ρ should be constrained in a certain “effective” range in DeepMAD.

Figure 3 gives more insights into the effectiveness hypothesis in Section 3.2. The architectures around $\rho = \{0.1, 0.5, 1.0\}$ are selected and grouped by ρ . When ρ is small (ρ is around 0.1), the network is effective in information propagation so we observe a strong correlation between network entropy and network accuracy. But these models

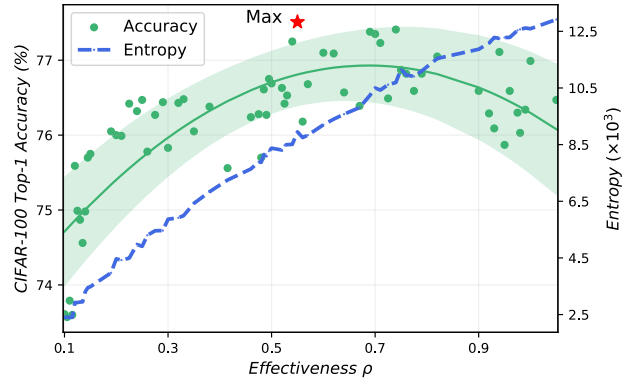


Figure 2. *Effectiveness* ρ v.s. top-1 accuracy and entropy of each generated model on CIFAR-100. The best model is marked by a star. The entropy increases with ρ monotonically but the model accuracy does not. The optimal $\rho^* \approx 0.5$.

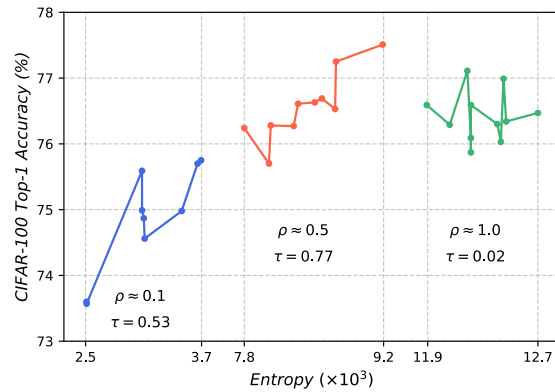


Figure 3. The architectures around $\rho = \{0.1, 0.5, 1.0\}$ are selected and grouped by ρ . Kendall coefficient τ [1] is used to measure the correlation.

are too shallow to obtain high performance. When ρ is too large ($\rho \approx 1.0$), the network approaches a chaos system therefore no clear correlation between network entropy and network accuracy. When ρ is around 0.5, the network can achieve the best performance and the correlation between the network entropy and network accuracy reaches 0.77.

5.4. DeepMAD for ResNet Family

ResNet family is one of the most popular and classic CNN models in deep learning. We use DeepMAD to re-design ResNet and show that the DeepMAD can generate much better ResNet models. The *effectiveness* ρ for those original ResNets is computed for easy comparison. First, we use DeepMAD to design a new architecture DeepMAD-R18 which has the same model size and FLOPs as ResNet-18. ρ is tuned in range $\{0.1, 0.3, 0.5, 0.7\}$ for DeepMAD-R18. $\rho = 0.3$ gives the best architecture. Then, $\rho = 0.3$

Model	# Param.	FLOPs	ρ	Acc. (%)
ResNet-18 [21]	11.7 M	1.8 G	0.01	70.9
ResNet-18†	11.7 M	1.8 G	0.01	72.2
DeepMAD-R18	11.7 M	1.8 G	0.1	76.9
DeepMAD-R18	11.7 M	1.8 G	0.3	77.7
DeepMAD-R18	11.7 M	1.8 G	0.5	77.5
DeepMAD-R18	11.7 M	1.8 G	0.7	75.7
ResNet-34 [21]	21.8 M	3.6 G	0.02	74.4
ResNet-34†	21.8 M	3.6 G	0.02	75.6
DeepMAD-R34	21.8 M	3.6 G	0.3	79.7
ResNet-50 [21]	25.6 M	4.1 G	0.09	77.4
ResNet-50†	25.6 M	4.1 G	0.09	79.3
DeepMAD-R50	25.6 M	4.1 G	0.3	80.6

Table 1. DeepMAD v.s. ResNet on ImageNet-1K, using ResNet building block. †: model trained by our pipeline. ρ is tuned for DeepMAD-R18. DeepMAD achieves consistent improvements compared with ResNet18/34/50 with the same Params and FLOPs.

is fixed in the design of DeepMAD-R34 and DeepMAD-R50 which align with ResNet-34 and Resnet-50 respectively. As shown in Table 1, compared to He’s original results, DeepMAD-R18 achieves 6.8% higher accuracy than ResNet-18 and is even comparable to ResNet-50. Besides, DeepMAD-R50 achieves 3.2% better accuracy than the ResNet-50. To ensure the fairness in comparison, the performances of ResNet family under the fair training setting are reported. With our training recipes, the accuracies of ResNet models improved around 1.5%. DeepMAD still outperforms the ResNet family by a large margin when both are trained fairly. The inferior performance of ResNet family can be explained by their small ρ which limits their model entropy. This phenomenon again validates our theory discussed in Section 4.3.

5.5. DeepMAD for Mobile CNNs

We use DeepMAD to design mobile CNN models for further exploration. Following previous works, MobileNet-V2 block with SE-block are used to build new models. ρ is tuned at EfficientNet-B0 scale in the range of {0.3, 0.5, 1.0, 1.5, 2.0} for DeepMAD-B0, and $\rho = 0.5$ achieves the best result. Then, we transfer the optimal ρ for DeepMAD-B0 to DeepMAD-MB. As shown in Table 2, the DeepMAD-B0 achieves 76.1% top-1 accuracy which is comparable with the EfficientNet-B0 (76.3%). It should be noted that EfficientNet-B0 is designed by brute-force grid search which takes around 3800 GPU days [66]. The performance of the DeepMAD-B0 is comparable to the EfficientNet-B0 by simply solving an MP problem on CPU in a few minutes. Aligned with MobileNet-V2 on Params and FLOPs, DeepMAD-MB achieves 72.3% top-1 accuracy which is 0.3% higher in accuracy.

Model	# Param.	FLOPs	ρ	Acc. (%)
EffNet-B0 [62]	5.3 M	390 M	0.6	76.3
DeepMAD-B0	5.3 M	390 M	0.3	74.3
DeepMAD-B0	5.3 M	390 M	0.5	76.1
DeepMAD-B0	5.3 M	390 M	1.0	75.9
DeepMAD-B0	5.3 M	390 M	1.5	75.7
DeepMAD-B0	5.3 M	390 M	2.0	74.9
MobileNet-V2 [23]	3.5 M	320 M	0.9	72.0
DeepMAD-MB	3.5 M	320 M	0.5	72.3

Table 2. DeepMAD under mobile setting. Top-1 accuracy on ImageNet-1K. ρ is tuned for DeepMAD-B0.

Model	Res	Params	FLOPs	Acc (%)
ResNet-50 [21]	224	26M	4.1G	77.4
DeiT-S [63]	224	22 M	4.6 G	79.8
PVT-Small [70]	224	25 M	3.8 G	79.8
Swin-T [40]	224	29 M	4.5 G	81.3
TNT-S [19]	224	24 M	5.2 G	81.3
T2T-ViT _t -14 [78]	224	22 M	6.1 G	81.7
ConvNeXt-T [41]	224	29 M	4.5 G	82.1
SLaK-T [39]	224	30 M	5.0 G	82.5
DeepMAD-29M	224	29 M	4.5 G	82.5
DeepMAD-29M*	288	29 M	4.5 G	82.8
ResNet-101 [21]	224	45M	7.8G	78.3
ResNet-152 [21]	224	60M	11.5G	79.2
PVT-Large [70]	224	61 M	9.8 G	81.7
T2T-ViT _t -19 [78]	224	39 M	9.8 G	82.2
T2T-ViT _t -24 [78]	224	64 M	15.0 G	82.6
TNT-B [19]	224	66 M	14.1 G	82.9
Swin-S [40]	224	50 M	8.7 G	83.0
ConvNeXt-S [41]	224	50 M	8.7 G	83.1
SLaK-S [39]	224	55 M	9.8 G	83.8
DeepMAD-50M	224	50 M	8.7 G	83.9
DeiT-B/16 [63]	224	87 M	17.6 G	81.8
RepLKNet-31B [16]	224	79 M	15.3 G	83.5
Swin-B [40]	224	88 M	15.4 G	83.5
ConvNeXt-B [41]	224	89 M	15.4 G	83.8
SLaK-B [39]	224	95 M	17.1 G	84.0
DeepMAD-89M	224	89 M	15.4 G	84.0

Table 3. DeepMAD v.s. SOTA ViT and CNN models on ImageNet-1K. $\rho = 0.5$ for all DeepMAD models. DeepMAD-29M*: uses 288x288 resolution while the Params and FLOPs keeps the same as DeepMAD-29M.

5.6. DeepMAD for SOTA

We use DeepMAD to design a SOTA CNN model for ImageNet-1K classification. The conventional MobileNet-V2 building block with SE module is used. This DeepMAD

network is aligned with Swin-Tiny [40] at 29M Params and 4.5G FLOPs therefore is labeled as DeepMAD-29M. As shown in Table 3, DeepMAD-29M outperforms or is comparable to SOTA ViT models as well as recent modern CNN models. DeepMAD-29M achieves 82.5%, which is 2.7% higher accuracy than DeiT-S [63] and 1.2% higher accuracy than the Swin-T [40]. Meanwhile, DeepMAD-29M is 0.4% higher than the ConvNeXt-T [41] which is inspired by the transformer architecture. DeepMAD also designs networks with larger resolution (288), DeepMAD-29M*, while keeping the FLOPs and Params not changed. DeepMAD-29M* reaches 82.8% accuracy and is comparable to Swin-S [40] and ConvNeXt-S [41] with nearly half of their FLOPs. Deep-MAD also achieves better performance on small and base level. Especially, DeepMAD-50M can achieve even better performance than ConvNeXt-B with nearly half of its scale. It proves only with the conventional convolutional layers as building blocks, Deep-MAD achieves comparable or better performance than ViT models.

5.7. Downstream Experiments

To demonstrate the transferability of models designed by DeepMAD, the models solved by DeepMAD play as the backbones on downstream tasks including object detection, semantic segmentation and action recognition.

Object Detection on MS COCO MS COCO is a widely used dataset in object detection. It has 143K images and 80 object categories. The experiments are evaluated on MS COCO [36] with the official training/testing splits. The results in Table 4 are evaluated on val-2017. We use two detection frameworks, FCOS [69] and GFLV2 [34], implemented by mmdetection [9]. The DeepMAD-R50 model plays as the backbone of these two detection frameworks. The models are initialized with pre-trained weights on ImageNet and trained for 2X (24 epochs). The multi-scale training trick is also used for the best performance. As shown in Table 4, DeepMAD-R50 achieves 40.0 AP with FCOS [69], which is 1.5 AP higher than ResNet-50. It also achieves 44.9 AP with GFLV2 [34], which is 1.0 AP higher than ResNet-50 again. The performance gain without introducing more Params and FLOPs proves the superiority of DeepMAD on network design.

Semantic segmentation on ADE20K ADE20K [81] dataset is broadly used in semantic segmentation tasks. It has 25k images and 150 semantic categories. The experiments are evaluated on ADE20K [81] with the official training/testing splits. The results in Table 5 are reported on testing part using mIoU. UperNet [72] in mmseg [11] is chosen as the segmentation framework. As shown in the first block of Table 5, DeepMAD-R50 achieves 45.6 mIoU on testing, which is 2.8 mIoU higher than ResNet-50, and

Backbone	# Param.	FLOPs	AP
FCOS			
ResNet-50	23.5 M	84.1 G	38.5
DeepMAD-R50	24.2 M	83.2 G	40.0
GFLV2			
ResNet-50	23.5 M	84.1 G	43.9
DeepMAD-R50	24.2 M	83.2 G	44.9

Table 4. DeepMAD for object detection and instance segmentation on MS COCO [36] with GFLV2 [34], FCOS [69], Mask R-CNN [20] and Cascade Mask R-CNN [7] frameworks. Backbones are pre-trained on ImageNet-1K. FLOPs and Params are counted for Backbone.

Backbone	# Param.	FLOPs	mIoU
ResNet-50	23.5 M	86.3 G	42.8
ResNet-101	42.5 M	164.3 G	44.8
DeepMAD-R50	24.2 M	85.2 G	45.6
Swin-T	27.5 M	95.8 G	45.8
ConvNeXt-T	27.8 M	93.2 G	46.7
DeepMAD-29M*	26.5 M	55.5 G	46.9

Table 5. DeepMAD for semantic segmentation on ADE20K [81]. All models are pre-trained on the ImageNet-1K and then fine-tuned using UperNet [72] framework. FLOPs and Params are counted for Backbone.

Backbone	# Param.	FLOPs	Acc. (%)
UCF-101			
ResNet-50	23.5 M	7.3 G	83.0
DeepMAD-R50	24.2 M	7.3 G	86.9
Kinetics-400			
ResNet-50	23.5 M	7.3 G	70.6
DeepMAD-R50	24.2 M	7.3 G	71.6

Table 6. DeepMAD for action recognition on UCF-101 [57] and Kinetics-400 [27] with the TSN [68] framework. Backbones are pre-trained on the ImageNet-1K. FLOPs and Params are counted for Backbone.

even 0.8 mIoU higher than ResNet-101. To compare to ViT and transformer-inspired models, DeepMAD-29M* is used as the backbone in UperNet. As shown in the last block of Table 5, DeepMAD-29M* achieves 46.9 mIoU on testing, which is 1.1 mIoU higher than Swin-T and 0.2 mIoU higher than ConvNeXt-T, with the same model size and less computation cost. It proves the advantage of CNN models designed by DeepMAD compared to transformer-based or transformer-inspired models.

Action recognition on UCF101 and Kinetics400 The UCF101 [57] dataset contains 13,320 video clips, covering 101 action classes. The Kinetics400 [27] dataset contains 400 human action classes, with more than 400 video clips for each action. They are both widely used in action recognition tasks. The results in Table 6 are reported on the testing part using top-1 accuracy. TSN [68] is adopted in mmaction [12] on UCF101 with split1 and Kinetics400 with official training/testing splits. As shown in Table 6, DeepMAD-R50 achieves 86.9% accuracy on UCF101 which is 3.9% higher than ResNet-50, and achieves 71.6% accuracy on Kinetics400 which is 1.0% higher than ResNet-50, with the same model size and computation cost. It shows that the models solved by DeepMAD can also be generalized to the recognition task on video datasets.

5.8. Ablation Study

In this section, we ablate important hyper-parameters and empirical guidelines in DeepMAD, with image classification on ImageNet-1K and object detection on COCO dataset. The **complexity comparison** is in Appendix F.

Ablation on entropy weights We generate networks with conventional convolution building blocks and different weight ratios α_i . The ratio α_5 is tuned in $\{1, 8, 16\}$ while the others are set to 1 as in [59]. As shown in Table 7, larger final stage weight can improve the performance on image classification task, while a smaller one can improve the performance on downstream task (object detection). For different tasks, additional improvements can be obtained by fine-tuning α_5 . However, this work uses $\alpha_5 = 8$ setting as the balance between the image classification task and object detection task. The experiments above have verified the advantage of DeepMAD on different tasks with global α_5 .

Ablation on the three empirical guidelines We generate networks using Mobilenet-V2 block with SE module and remove one of the three empirical guidelines discussed in Section 4.2 at each time to explore their influence. As shown in Table 8, removing any one of the three guidelines will degrade the performance of the model. Particularly, the third guideline is the most critical one for DeepMAD.

6. Limitations

As no research is perfect, DeepMAD has several limitations as well. First, three empirical guidelines discussed in Section 4.2 do not have strong theoretical foundation. Hopefully, they can be removed or replaced in the future. Second, DeepMAD has several hyper-parameters to tune, such as $\{\alpha_i\}$ and $\{\beta, \rho\}$. Third, DeepMAD focuses on conventional CNN layers at this stage while there are many

Model	# Param.	FLOPs	α_5	Acc. (%)
DeepMAD-R18	11.7M	1.8G	1	76.7
DeepMAD-R18	11.7M	1.8G	8	77.7
DeepMAD-R18	11.7M	1.8G	16	78.7
Backbone	# Param.	FLOPs	α_5	AP (%)
DeepMAD-R18	9.8M	37.0G	1	35.1
DeepMAD-R18	11.0M	36.8G	8	34.6
DeepMAD-R18	11.1M	36.8G	16	34.1

Table 7. The performance on ImageNet-1k and COCO dataset of DeepMAD-R18 with different final stage weight α_5 . $\alpha_5 = 8$ balances the good performance between classification and object detection and is adopted in this work.

Guideline1	Guideline2	Guideline3	Acc. (%)
×	✓	✓	73.5
✓	×	✓	73.5
✓	✓	×	73.1
✓	✓	✓	73.7

Table 8. Top-1 accuracies on ImageNet-1K of DeepMAD-B0 with different combination of guidelines. The model designed with all three guidelines achieves the best results. All models are trained for 120 epochs.

more powerful and more modern building blocks such as transformers. It is potentially possible to generalize DeepMAD to these building blocks as well in future works.

7. Conclusion

We propose a pure mathematical framework DeepMAD for designing high-performance convolutional neural networks. The key idea of the DeepMAD is to maximize the network entropy while keeping network effectiveness bounded by a small constant. We show that DeepMAD can design SOTA CNN models that are comparable to or even better than ViT models and modern CNN models. To demonstrate the power of DeepMAD, we only use conventional convolutional building blocks, like ResNet block, and depth-wise convolution in MobileNet-V2. Without bells and whistles, DeepMAD achieves competitive performance using these old-school building blocks. This encouraging result implies that the full potential of the conventional CNN models has not been fully released due to the previous sub-optimal design. Hope this work can attract more research attention to theoretical deep learning in the future.

8. Acknowledgment

This work was supported by Alibaba Research Intern Program and National Science Foundation CCF-1919117.

References

- [1] Hervé Abdi. The kendall rank correlation coefficient. **5**
- [2] Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018. **2**
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. **1**
- [4] Thomas Bäck. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, Inc., USA, 1996. **4**
- [5] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *ICLR*, 2020. **1**
- [6] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. **14**
- [7] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. *CVPR*, pages 6154–6162, 2018. **7**
- [8] Kwan Ho Ryan Chan, Yaodong Yu, Chong You, Haozhi Qi, John Wright, and Yi Ma. Redunet: A white-box deep network from the principle of maximizing rate reduction. *arXiv preprint arXiv:2105.10446*, 2021. **2, 3**
- [9] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. **7**
- [10] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *ICLR*, 2021. **2**
- [11] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. **7**
- [12] MMAction2 Contributors. Openmmlab’s next generation video understanding toolbox and benchmark. <https://github.com/open-mmlab/mmaction2>, 2020. **8**
- [13] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999. **12**
- [14] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. **5, 15**
- [15] 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. **4**
- [16] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *CVPR*, pages 11963–11975, 2022. **1, 2, 6**
- [17] Peiyan Dong, Mengshu Sun, Alec Lu, Yanyue Xie, Kenneth Liu, Zhenglun Kong, Xin Meng, Zhengang Li, Xue Lin, Zhenman Fang, et al. Heatvit: Hardware-efficient adaptive token pruning for vision transformers. *arXiv preprint arXiv:2211.08110*, 2022. **1**
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. **1, 2, 5**
- [19] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *NeurIPS*, 34:15908–15919, 2021. **6**
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:386–397, 2020. **7**
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. **2, 5, 6, 14**
- [22] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *ICCV*, 2021. **1, 2**
- [23] 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. **1, 2, 5, 6, 14**
- [24] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018. **5**
- [25] Yeran Idelbayev. Proper ResNet implementation for CIFAR10/CIFAR100 in PyTorch. https://github.com/akamaster/pytorch_resnet_cifar10. **5**
- [26] E. T. Jaynes. Information theory and statistical mechanics. In *Phys. Rev.* 106, 620, 1957. **2, 3**
- [27] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. **5, 7, 8**
- [28] Zhenglun Kong, Peiyan Dong, Xiaolong Ma, Xin Meng, Wei Niu, Mengshu Sun, Bin Ren, Minghai Qin, Hao Tang, and Yanzhi Wang. Spvit: Enabling faster vision transformers via soft token pruning. *arXiv preprint arXiv:2112.13890*, 2021. **1**
- [29] Zhenglun Kong, Haoyu Ma, Geng Yuan, Mengshu Sun, Yanyue Xie, Peiyan Dong, Xin Meng, Xuan Shen, Hao Tang, Minghai Qin, et al. Peeling the onion: Hierarchical reduction of data redundancy for efficient vision transformer training. *arXiv preprint arXiv:2211.10801*, 2022. **1**
- [30] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research). **4**
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. **1**

- [32] Solomon Kullback. Information theory and statistics. 1997. [2, 3](#)
- [33] Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, and Xiaojun Chang. Blockwisely supervised neural architecture search with knowledge distillation. In *CVPR*, pages 1989–1998, 2020. [2](#)
- [34] Xiang Li, Wenhai Wang, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. *arXiv preprint arXiv:2011.12885*, 2020. [7](#)
- [35] Ming Lin, Pichao Wang, Zhenhong Sun, Hesen Chen, Xiuyu Sun, Qi Qian, Hao Li, and Rong Jin. Zen-nas: A zero-shot nas for high-performance deep image recognition. In *ICCV*, 2021. [2](#)
- [36] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. [5, 7](#)
- [37] 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. [2](#)
- [38] 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. [14](#)
- [39] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Xuxi Chen, Qiao Xiao, Boqian Wu, Mykola Pechenizkiy, Decebal Mocanu, and Zhangyang Wang. More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. *arXiv preprint arXiv:2207.03620*, 2022. [6, 13](#)
- [40] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. [1, 2, 5, 6, 7, 13](#)
- [41] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CVPR*, 2022. [1, 2, 6, 7, 13](#)
- [42] Michel Loeve. *Probability theory*. Courier Dover Publications, 2017. [12](#)
- [43] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [5](#)
- [44] MathWorks. *MATLAB:R2022b*. The MathWorks Inc., Natick, Massachusetts, 2022. [4](#)
- [45] Sachin Mehta and Mohammad Rastegari. Mobilevit: Lightweight, general-purpose, and mobile-friendly vision transformer. *arXiv preprint arXiv:2110.02178*, 2021. [1](#)
- [46] Alexander McFarlane Mood. Introduction to the theory of statistics. 1950. [12](#)
- [47] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *ICML*, pages 4095–4104. PMLR, 2018. [5](#)
- [48] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, pages 10428–10436, 2020. [2](#)
- [49] 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. [14](#)
- [50] Daniel A. Roberts, Sho Yaida, and Boris Hanin. *The Principles of Deep Learning Theory*. Cambridge University Press, 2022. <https://deeplearningtheory.com>. [2, 3](#)
- [51] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018. [1](#)
- [52] Andrew Michael Saxe, Yamini Bansal, Joel Dapello, Madhu Advani, Artemy Kolchinsky, Brendan Daniel Tracey, and David Daniel Cox. On the information bottleneck theory of deep learning. In *ICLR*, 2018. [3](#)
- [53] Christian Szegedy Sergey Ioffe. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *arXiv:1502.03167*, 2015. [2](#)
- [54] Xuan Shen, Zhenglun Kong, Minghai Qin, Peiyan Dong, Geng Yuan, Xin Meng, Hao Tang, Xiaolong Ma, and Yanzhi Wang. The lottery ticket hypothesis for vision transformers. *arXiv preprint arXiv:2211.01484*, 2022. [1](#)
- [55] Yao Shu, Shaofeng Cai, Zhongxiang Dai, Beng Chin Ooi, and Bryan Kian Hsiang Low. Nasi: Label-and data-agnostic neural architecture search at initialization. *arXiv preprint arXiv:2109.00817*, 2021. [2](#)
- [56] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [2](#)
- [57] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [5, 7, 8](#)
- [58] Zhenhong Sun, Ce Ge, Junyan Wang, Ming Lin, Hesen Chen, Hao Li, and Xiuyu Sun. Entropy-driven mixed-precision quantization for deep network design. In *Advances in Neural Information Processing Systems*. [2](#)
- [59] Zhenhong Sun, Ming Lin, Xiuyu Sun, Zhiyu Tan, Hao Li, and Rong Jin. Mae-det: Revisiting maximum entropy principle in zero-shot nas for efficient object detection. In *International Conference on Machine Learning*, pages 20810–20826. PMLR, 2022. [2, 3, 4, 8](#)
- [60] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. [5](#)
- [61] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*, pages 2820–2828, 2019. [2](#)
- [62] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019. [1, 2, 5, 6](#)
- [63] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *ICML*, volume 139, pages 10347–10357, July 2021. [1, 2, 5, 6, 7](#)

- [64] Mac E Van Valkenburg. *Reference data for engineers: radio, electronics, computers and communications*. Newnes, 2001. [12](#)
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NerulIPS*, volume 30. Curran Associates, Inc., 2017. [1](#)
- [66] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuan-dong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, and Joseph E. Gonzalez. Fb-netv2: Differentiable neural architecture search for spatial and channel dimensions. In *CVPR*, pages 12962–12971, 2020. [6](#)
- [67] Junyan Wang, Zhenhong Sun, Yichen Qian, Dong Gong, Xiuyu Sun, Ming Lin, Maurice Pagnucco, and Yang Song. Maximizing spatio-temporal entropy of deep 3d CNNs for efficient video recognition. In *The Eleventh International Conference on Learning Representations, 2023*. [2](#)
- [68] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36. Springer, 2016. [7](#), [8](#)
- [69] Ning Wang, Yang Gao, Hao Chen, Peng Wang, Zhi Tian, Chunhua Shen, and Yanning Zhang. Nas-fcos: Fast neural architecture search for object detection. In *CVPR*, pages 11943–11951, 2020. [7](#)
- [70] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021. [6](#)
- [71] Ross Wightman, Hugo Touvron, and Herve Jegou. Resnet strikes back: An improved training procedure in timm. In *NerulIPS*, 2021. [5](#)
- [72] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, pages 418–434, 2018. [7](#)
- [73] Junyuan Xie, Tong He, Zhi Zhang, Hang Zhang, and Jerry Zhang. Bag of tricks for image classification with convolutional neural networks. In *CVPR*, 2019. [5](#)
- [74] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018. [2](#)
- [75] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018. [14](#)
- [76] Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. Metaformer is actually what you need for vision. In *CVPR*, pages 10819–10829, 2022. [1](#), [2](#)
- [77] Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. *NerulIPS*, 33:9422–9434, 2020. [2](#), [3](#)
- [78] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, pages 558–567, 2021. [6](#), [13](#)
- [79] Miao Zhang, Huiqi Li, Shirui Pan, Xiaojun Chang, and Steven Su. Overcoming multi-model forgetting in one-shot nas with diversity maximization. In *CVPR*, pages 7806–7815, 2020. [5](#)
- [80] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, volume 34, pages 13001–13008, 2020. [5](#)
- [81] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, pages 5122–5130, 2017. [5](#), [7](#)
- [82] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *CVPR*, pages 8697–8710, 2018. [14](#)