

InternImage: Exploring Large-Scale Vision Foundation Models with Deformable Convolutions

Wenhai Wang^{1*}, Jifeng Dai^{2,1*}, Zhe Chen^{3,1*}, Zhenhang Huang^{1*}, Zhiqi Li^{3,1*}, Xizhou Zhu^{4*}, Xiaowei Hu¹, Tong Lu³, Lewei Lu⁴, Hongsheng Li⁵, Xiaogang Wang^{4,5}, Yu Qiao^{1✉}

¹Shanghai AI Laboratory ²Tsinghua University

³Nanjing University ⁴SenseTime Research ⁵The Chinese University of Hong Kong

<https://github.com/OpenGVLab/InternImage>

Abstract

Compared to the great progress of large-scale vision transformers (ViTs) in recent years, large-scale models based on convolutional neural networks (CNNs) are still in an early state. This work presents a new large-scale CNN-based foundation model, termed InternImage, which can obtain the gain from increasing parameters and training data like ViTs. Different from the recent CNNs that focus on large dense kernels, InternImage takes deformable convolution as the core operator, so that our model not only has the large effective receptive field required for downstream tasks such as detection and segmentation, but also has the adaptive spatial aggregation conditioned by input and task information. As a result, the proposed InternImage reduces the strict inductive bias of traditional CNNs and makes it possible to learn stronger and more robust patterns with large-scale parameters from massive data like ViTs. The effectiveness of our model is proven on challenging benchmarks including ImageNet, COCO, and ADE20K. It is worth mentioning that InternImage-H achieved a new record 65.4 mAP on COCO test-dev and 62.9 mIoU on ADE20K, outperforming current leading CNNs and ViTs.

1. Introduction

With the remarkable success of transformers in large-scale language models [3–8], vision transformers (ViTs) [2, 9–15] have also swept the computer vision field and are becoming the primary choice for the research and practice of large-scale vision foundation models. Some pioneers [16–20] have made attempts to extend ViTs to very large models with over a billion parameters, beating convolutional neural networks (CNNs) and significantly pushing the performance bound for a wide range of computer vision tasks, including basic classification, detection, and segmen-

* equal contribution, ✉ corresponding author (qiaoyu@pjlab.org.cn)

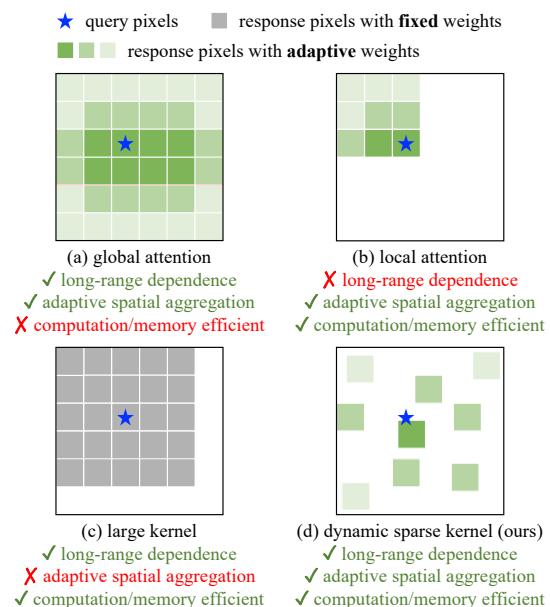


Figure 1. **Comparisons of different core operators.** (a) shows the global aggregation of multi-head self-attention (MHSA) [1], whose computational and memory costs are expensive in downstream tasks that require high-resolution inputs. (b) limits the range of MHSA into a local window [2] to reduce the cost. (c) is a depth-wise convolution with very large kernels to model long-range dependencies. (d) is a deformable convolution, which shares similar favorable properties with MHSA and is efficient enough for large-scale models. We start from it to build a large-scale CNN.

tation. While these results suggest that CNNs are inferior to ViTs in the era of massive parameters and data, we argue that *CNN-based foundation models can also achieve comparable or even better performance than ViTs when equipped with similar operator-/architecture-level designs, scaling-up parameters, and massive data.*

To bridge the gap between CNNs and ViTs, we first summarize their differences from two aspects: (1) From the operator level [9, 21, 22], the multi-head self-attention

(MHSA) of ViTs has long-range dependencies and adaptive spatial aggregation (see Fig. 1(a)). Benefiting from the flexible MHSA, ViTs can learn more powerful and robust representations than CNNs from massive data. (2) From the architecture view [9, 22, 23], besides MHSA, ViTs contain a series of advanced components that are not included in standard CNNs, such as Layer Normalization (LN) [24], feed-forward network (FFN) [1], GELU [25], etc. Although recent works [21, 22] have made meaningful attempts to introduce long-range dependencies into CNNs by using dense convolutions with very large kernels (*e.g.*, 31×31) as shown in Fig. 1 (c), there is still a considerable gap with the state-of-the-art large-scale ViTs [16, 18–20, 26] in terms of performance and model scale.

In this work, we concentrate on designing a CNN-based foundation model that can efficiently extend to large-scale parameters and data. Specifically, we start with a flexible convolution variant—deformable convolution (DCN) [27, 28]. By combining it with a series of tailored block-level and architecture-level designs similar to transformers, we design a brand-new convolutional backbone network, termed *InternImage*. As shown in Fig. 1, different from recently improved CNNs with very large kernels such as 31×31 [22], the core operator of InternImage is a dynamic sparse convolution with a common window size of 3×3 , (1) whose sampling offsets are flexible to dynamically learn appropriate receptive fields (can be long- or short-range) from given data; (2) the sampling offsets and modulation scalars are adaptively adjusted according to the input data, which can achieve adaptive spatial aggregation like ViTs, reducing the over-inductive bias of regular convolutions; and (3) the convolution window is a common 3×3 , avoiding the optimization problems and expensive costs caused by large dense kernels [22, 29].

With the aforementioned designs, the proposed InternImage can efficiently scale to large parameter sizes and learn stronger representations from large-scale training data, achieving comparable or even better performance to large-scale ViTs [2, 11, 19] on a wide range of vision tasks. In summary, our main contributions are as follows:

(1) We present a new large-scale CNN-based foundation model—InternImage. To our best knowledge, it is the first CNN that effectively scales to over 1 billion parameters and 400 million training images and achieves comparable or even better performance than state-of-the-art ViTs, showing that convolutional models are also a worth-exploring direction for large-scale model research.

(2) We successfully scale CNNs to large-scale settings by introducing long-range dependencies and adaptive spatial aggregation using an improved 3×3 DCN operator, and explore the tailored basic block, stacking rules, and scaling strategies centered on the operator. These designs make effective use of the operator, enabling our models to obtain

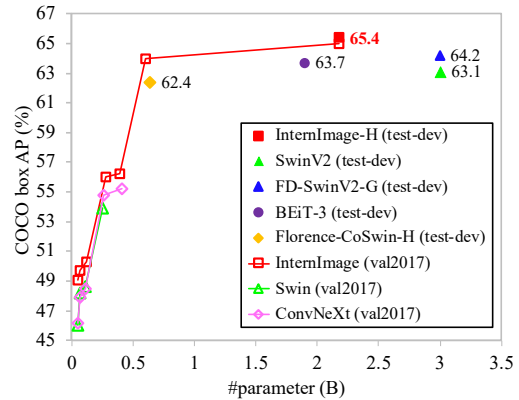


Figure 2. **Performance comparison on COCO of different backbones.** The proposed InternImage-H achieves a new record 65.4 box AP on COCO test-dev, significantly outperforming state-of-the-art CNNs and large-scale ViTs.

the gains from large-scale parameters and data.

(3) We evaluate the proposed model on representative vision tasks including image classification, object detection, instance and semantic segmentation, and compared it with state-of-the-art CNNs and large-scale ViTs by scaling the model size ranging from 30 million to 1 billion, the data ranging from 1 million to 400 million. Specifically, our model with different parameter sizes can consistently outperform prior arts on ImageNet [30]. InternImage-B achieves 84.9% top-1 accuracy trained only on the ImageNet-1K dataset, outperforming CNN-based counterparts [21, 22] by at least 1.1 points. With large-scale parameters (*i.e.*, 1 billion) and training data (*i.e.*, 427 million), the top-1 accuracy of InternImage-H is further boosted to 89.6%, which is close to well-engineering ViTs [2, 19] and hybrid-ViTs [20]. In addition, on COCO [31], a challenging downstream benchmark, our best model InternImage-H achieves state-of-the-art 65.4% box mAP with 2.18 billion parameters, 2.3 points higher than SwinV2-G [16] (65.4 vs. 63.1) with 27% fewer parameters as shown in Fig. 2.

2. Related Work

Vision foundation models. Convolutional neural networks (CNNs) became the mainstream for visual recognition after the large-scale dataset and computation resources were available. Straining from AlexNet [32], lots of deeper and more effective neural network architectures have been proposed, such as VGG [33], GoogleNet [34], ResNet [35], ResNeXt [36], EfficientNet [37, 38], etc. In addition to the architectural design, more sophisticated convolution operations such as depth-wise convolution [39] and deformable convolution [27, 28] are formulated. By considering the advanced designs of transformers, modern CNNs showed promising performance on the vision tasks by discovering better components in macro/micro designs and intro-

ducing improved convolutions with long-range dependencies [21, 22, 40] or dynamic weights [41].

In recent years, a new line of vision foundation models focuses on transformer-based architecture. ViT [9] is the most representative model, which achieves great success in vision tasks thanks to global receptive fields and dynamic spatial aggregation. However, global attention in ViT suffers from expensive computational/memory complexity, especially on large feature maps, which limits its application in downstream tasks. To address this problem, PVT [10, 11] and Linformer [42] performed global attention on the downsampled key and value maps, DAT [43] employed deformable attention to sparsely sample information from value maps, while HaloNet [44] and Swin transformer [2] developed local attention mechanisms and used haloing and shift operations to transfer information among adjacent local regions.

Large-scale models. Scaling up models is an important strategy to improve feature representation quality, which has been well-studied in the natural language processing (NLP) domain [45]. Inspired by the success in the NLP field, Zhai *et al.* [19] first extended ViT to 2 billion parameters. Liu *et al.* [16] enlarged the hierarchical-structure Swin transformer to a deeper and wider model with 3 billion parameters. Some researchers developed large-scale hybrid ViTs [20, 46] by combining the advantages of ViTs and CNNs at different levels. Recently, BEiT-3 [17] further explored stronger representations based on ViT with large-scale parameters using multimodal pre-training. These methods significantly raise the upper bound of basic vision tasks. However, research on CNN-based large-scale models has lagged behind transformer-based architectures in terms of the total number of parameters and performance. Although newly-proposed CNNs [21, 22, 40, 47] introduce long-range dependencies by using convolutions with very large kernels or recursive gated kernels, there is still a considerable gap with state-of-the-art ViTs. In this work, we aim to develop a CNN-based foundation model that can extend efficiently to a large scale comparable to ViT.

3. Proposed Method

To design a large-scale CNN-based foundation model, we start with a flexible convolution variant, namely deformable convolution v2 (DCNv2) [28] and make some tune-ups based on it to better suit the requirements of large-scale foundation models. Then, we build the basic block by combining the tuned convolution operator with advanced block designs used in modern backbones [16, 19]. Finally, we explore the stacking and scaling principles of DCN-based blocks to build a large-scale convolutional model that can learn strong representations from massive data.

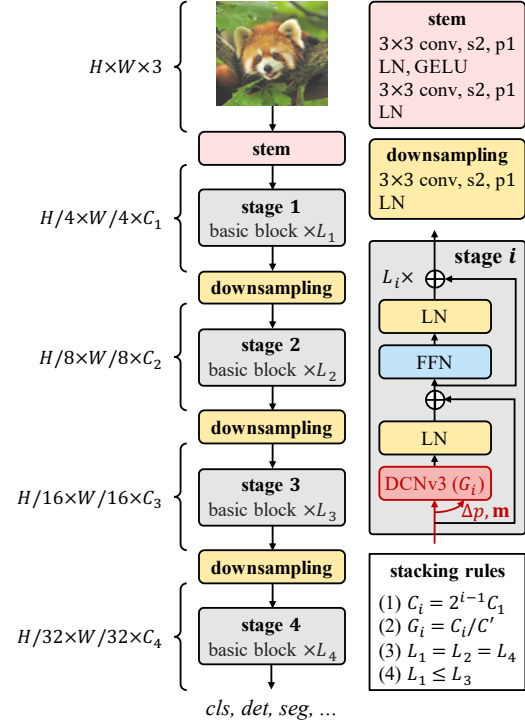


Figure 3. **Overall Architecture of InternImage**, where the core operator is DCNv3, and the basic block composes of layer normalization (LN) [24] and feed-forward network (FFN) [1] as transformers, the stem and downsampling layers follows conventional CNN’s designs, where “s2” and “p1” mean stride 2 and padding 1, respectively. Constrained by the stacking rules, only 4 hyper-parameters (C_1, C', L_1, L_3) can decide a model variant.

3.1. Deformable Convolution v3

Convolution vs. MHSA. Previous works [21, 22, 48] have extensively discussed the differences between CNNs and ViTs. Before deciding on the core operator of InternImage, we first summarize the main differences between regular convolution and MHSA.

(1) *Long-range dependencies.* Although it has long been recognized that models with large effective receptive fields (long-range dependencies) usually perform better on downstream vision tasks [49–51], the de-facto effective receptive field of CNNs [33, 35] stacked by 3×3 regular convolutions is relatively small. Even with very deep models, the CNN-based model still cannot acquire long-range dependencies like ViTs, which limits its performance.

(2) *Adaptive spatial aggregation.* Compared to MHSA whose weights are dynamically conditioned by the input, regular convolution [52] is an operator with static weights and strong inductive biases such as 2D locality, neighborhood structure, translation equivalence, etc. With the highly-inductive properties, models composed by regular convolutions might converge faster and require less training data than ViTs, but it also restricts CNNs from learn-

ing more general and robust patterns from web-scale data. More robustness experiments are detailed in the supplementary material.

Revisiting DCNv2. A straightforward way to bridge the gap between convolution and MHSA is to introduce long-range dependencies and adaptive spatial aggregation into regular convolutions. Let us start with DCNv2 [28], which is a general variant of regular convolution. Given an input $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ and current pixel p_0 , DCNv2 can be formulated as:

$$\mathbf{y}(p_0) = \sum_{k=1}^K \mathbf{w}_k \mathbf{m}_k \mathbf{x}(p_0 + p_k + \Delta p_k), \quad (1)$$

where K represents the total number of sampling points, and k enumerates the sampling point. $\mathbf{w}_k \in \mathbb{R}^{C \times C}$ denotes the projection weights of the k -th sampling point, and $\mathbf{m}_k \in \mathbb{R}$ represents the modulation scalar of the k -th sampling point, which is normalized by sigmoid function. p_k denotes the k -th location of the pre-defined grid sampling $\{(-1, -1), (-1, 0), \dots, (0, +1), \dots, (+1, +1)\}$ as in regular convolutions, and Δp_k is the offset corresponding to the k -th grid sampling location. We see from the equation that (1) for long-range dependencies, the sampling offset Δp_k is flexible and able to interact with short- or long-range features; and (2) for adaptive spatial aggregation, both the sampling offset Δp_k and modulation scalar \mathbf{m}_k are learnable and conditioned by input \mathbf{x} . So it can be found that *DCNv2 shares similar favorable properties with MHSA*, which motivated us to develop large-scale CNN-based foundation models on the basis of this operator.

Extending DCNv2 for Vision Foundation Models. In common practice, DCNv2 is usually used as an extension to regular convolutions, loading pre-trained weights of regular convolutions and fine-tuning for better performance, which is not exactly suitable for large-scale vision foundation models that need to be trained from scratch. In this work, to address this problem, we extend DCNv2 from aspects as follows:

(1) *Sharing weights among convolutional neurons.* Similar to regular convolution, different convolutional neurons¹ in original DCNv2 have independent linear projection weights, and thus its parameter and memory complexity are linear with the total number of sampling points, which significantly limits the efficiency of the model, especially in large-scale models. To remedy this problem, we borrow the idea from the separable convolution [53] and detach the original convolution weights \mathbf{w}_k into depth-wise and point-wise parts, where the depth-wise part is responsible by the original location-aware modulation scalar \mathbf{m}_k , and the point-wise part is the shared projection weights \mathbf{w} among sampling points.

¹A 3×3 regular convolution has 9 linear projection neurons.

(2) *Introducing multi-group mechanism.* The multi-group (head) design first appeared in group convolution [32], and it is widely used in MHSA [1] of transformers and works with adaptive spatial aggregation to effectively learn richer information from different representation subspaces at different locations. Inspired by this, we split the spatial aggregation process into G groups, each of which has individual sampling offsets Δp_{gk} and modulation scale \mathbf{m}_{gk} , and thus different groups on a single convolution layer can have different spatial aggregation patterns, resulting in stronger features for downstream tasks.

(3) *Normalizing modulation scalars along sampling points.* The modulation scalars in the original DCNv2 are element-wise normalized by the sigmoid function. Therefore, each modulation scalar is in the range $[0, 1]$, and the sum of the modulation scalars of all sample points is not stable and varies from 0 to K . This leads to unstable gradients in DCNv2 layers when training with large-scale parameters and data. To alleviate the instability issues, we change element-wise sigmoid normalization to softmax normalization along sample points. In this way, the sum of the modulation scalars is constrained to 1, which makes the training process of models at different scales more stable.

Combining the mentioned modifications, the extended DCNv2, marked as DCNv3, can be formulated as Eqn. (2).

$$\mathbf{y}(p_0) = \sum_{g=1}^G \sum_{k=1}^K \mathbf{w}_g \mathbf{m}_{gk} \mathbf{x}_g(p_0 + p_k + \Delta p_{gk}), \quad (2)$$

where G denotes the total number of aggregation groups. For the g -th group, $\mathbf{w}_g \in \mathbb{R}^{C \times C'}$, $\mathbf{m}_{gk} \in \mathbb{R}$ denote the location-irrelevant projection weights of the group, where $C' = C/G$ represents the group dimension. $\mathbf{m}_{gk} \in \mathbb{R}$ denotes the modulation scalar of the k -th sampling point in the g -th group, normalized by the softmax function along the dimension K . $\mathbf{x}_g \in \mathbb{R}^{C' \times H \times W}$ represents the sliced input feature map. Δp_{gk} is the offset corresponding to the grid sampling location p_k in the g -th group.

In general, DCNv3, as an extension of the DCN series, enjoys three merits as follows: (1) This operator made up for the deficiencies of regular convolution in terms of long-range dependencies and adaptive spatial aggregation; (2) Compared with attention-based operators such as common MHSA and closely-related deformable attention [43, 54], this operator inherits the inductive bias of convolution, making our model more efficient with fewer training data and shorter training time; (3) This operator is based on sparse sampling, which is more computational and memory efficient than previous methods such as MHSA [1] and re-parameterizing large kernel [22]. In addition, due to the sparse sampling, DCNv3 only needs a 3×3 kernel to learn long-range dependencies, which is easier to be optimized and avoids extra auxiliary techniques such as re-parameterizing [22] used in large kernels.

3.2. InternImage Model

Using DCNv3 as the core operator brings a new problem: *how to build a model that can make effective use of the core operator?* In this section, we first present the details of the basic block and other integral layers of our model, and then we construct a new CNN-based foundation model termed InternImage, by exploring a tailored stacking strategy for these basic blocks. Finally, we study scaling-up rules for the proposed model to obtain the gain from increasing parameters.

Basic block. Unlike the widely used bottlenecks in traditional CNNs [35], the design of our basic block is closer to ViTs, which is equipped with more advanced components including LN [24], feed-forward networks (FFN) [1], and GELU [25]. This design is proved to be efficient [2, 10, 11, 21, 22] in various vision tasks. The details of our basic block are illustrated in Fig. 3, where the core operator is DCNv3, and the sampling offsets and modulation scales are predicted by passing input feature x through a separable convolution (a 3×3 depth-wise convolution followed by a linear projection). For other components, we use the post-normalization setting [55] by default and follow the same design as that of the plain transformer [1, 9].

Stem & downsampling layers. To obtain hierarchical feature maps, we use convolutional stem and downsampling layers to resize the feature maps to different scales. As shown in Fig. 3, the stem layer is placed before the first stage to reduce the input resolution by 4 times. It consists of two convolutions, two LN layers, and one GELU layer, where the kernel size of the two convolutions is 3, the stride is 2, the padding is 1, and the output channel of the first convolution is half of the second one. Similarly, the downsampling layer is made up of a 3×3 convolution with a stride of 2 and a padding of 1, followed by one LN layer. It sits between the two stages and is used to downsample the input feature map by 2 times.

Stacking rules. To clarify the block-stacking process, we first list the hyper-parameters of InternImage as follows:

- C_i : the channel number of the i -th stage;
- G_i : the group number of the DCNv3 in the i -th stage;
- L_i : the number of basic blocks in the i -th stage.

Since our model has 4 stages, a variant is decided by 12 hyper-parameters, whose search space is too large to exhaustively enumerate and find the best variant. To reduce the search space, we summarize the design experiences of prior arts [2, 21, 35] into 4 rules as shown in Fig. 3, where the first rule makes the channel numbers of the last three stages determined by the channel number C_1 of the first stage, and the second rule lets the group number correspond to the channel number of stages. For the number of stacked blocks in different stages, we simplify the stacking pattern to “AABA”, which means the block number of stage 1, 2, and 4 are the same, and are not greater than that of the stage

model name	C_1	C'	$L_{1,2,3,4}$	#params
InternImage-T (origin)	64	16	4, 4, 18, 4	30M
InternImage-S	80	16	4, 4, 21, 4	50M
InternImage-B	112	16	4, 4, 21, 4	97M
InternImage-L	160	16	5, 5, 22, 5	223M
InternImage-XL	192	16	5, 5, 24, 5	335M
InternImage-H	320	32	6, 6, 32, 6	1.08B

Table 1. **Hyper-parameters for models of different scales.** InternImage-T is the origin model, and -S/B/L/XL/H are scaled up from -T. “#params” denotes the number of parameters.

3 as illustrated in the last two rules. With these rules, a InternImage variant can be defined by using only 4 hyper-parameters (C_1, C', L_1, L_3).

Let us choose a model with 30 million parameters as the origin and discretize C_1 to $\{48, 64, 80\}$, L_1 to $\{1, 2, 3, 4, 5\}$, and C' to $\{16, 32\}$. In this way, the original huge search space is reduced to 30, and we can find the best model from the 30 variants by training and evaluating them in ImageNet [30]. In practice, we use the best hyper-parameter setting (64, 16, 4, 18) to define the origin model and scale it to different scales.

Scaling rules. Based on the optimal origin model under the aforementioned constraints, we further explore the parameter scaling rules inspired by [37]. Specifically, we consider two scaling dimensions: depth D (i.e., $3L_1 + L_3$) and width C_1 , and scale the two dimensions using α, β and a composite factor ϕ . The scaling rules can be written as: $D' = \alpha^\phi D$ and $C'_1 = \beta^\phi C_1$,

where $\alpha \geq 1, \beta \geq 1$, and $\alpha\beta^{1.99} \approx 2$. Here, 1.99 is specific for InternImage and calculated by doubling the model width and keeping the depth constant. We experimentally find out that the best scaling setting is $\alpha = 1.09$ and $\beta = 1.36$, and then we base on it to construct InternImage variants with different parameter scales, namely InternImage-T/S/B/L/XL, whose complexity is similar to those of ConvNeXt [21]. To further test the capability, we built a larger InternImage-H with 1 billion parameters, and to accommodate very large model widths, we also change the group dimension C' to 32. The configurations are summarized in Table 1.

4. Experiment

We analyze and compare InternImage with the leading CNNs and ViTs on representative vision tasks including image classification, object detection, instance and semantic segmentation. Besides the experiments in the main paper, due to space constraints, more experimental setups and ablation studies are presented in the supplementary materials.

4.1. Image Classification

Settings. We evaluate the classification performance of InternImage on ImageNet [30]. For fair comparisons, following common practices [2, 10, 21, 56], InternImage-T/S/B are trained on ImageNet-1K (~ 1.3 million) for 300 epochs,

method	type	scale	#params	#FLOPs	acc (%)
Swin-T [2]	T	224 ²	29M	5G	81.3
CoAtNet-0 [20]	T	224 ²	25M	4G	81.6
PVTv2-B2 [11]	T	224 ²	25M	4G	82.0
DeiT III-S [62]	T	224 ²	22M	5G	81.4
SwinV2-T/8 [16]	T	256 ²	28M	6G	81.8
ConvNeXt-T [21]	C	224 ²	29M	5G	82.1
InternImage-T (ours)	C	224²	30M	5G	83.5
Swin-S [2]	T	224 ²	50M	9G	83.0
CoAtNet-1 [20]	T	224 ²	42M	8G	83.3
PVTv2-B4 [11]	T	224 ²	63M	10G	83.6
SwinV2-S/8 [16]	T	256 ²	50M	12G	83.7
ConvNeXt-S [21]	C	224 ²	50M	9G	83.1
InternImage-S (ours)	C	224²	50M	8G	84.2
Swin-B [2]	T	224 ²	88M	15G	83.5
CoAtNet-2 [20]	T	224 ²	75M	16G	84.1
PVTv2-B5 [11]	T	224 ²	82M	12G	83.8
DeiT III-B [62]	T	224 ²	87M	18G	83.8
SwinV2-B/8 [16]	T	256 ²	88M	20G	84.2
RepLkNet-31B [22]	C	224 ²	79M	15G	83.5
ConvNeXt-B [21]	C	224 ²	88M	15G	83.8
InternImage-B (ours)	C	224²	97M	16G	84.9
Swin-L [‡] [2]	T	384 ²	197M	104G	87.3
CoAtNet-4 [‡] [20]	T	384 ²	275M	190G	87.9
DeiT III-L [‡] [62]	T	384 ²	304M	191G	87.7
SwinV2-L/24 [‡] [16]	T	384 ²	197M	115G	87.6
RepLkNet-31L [‡] [22]	C	384 ²	172M	96G	86.6
ConvNeXt-L [‡] [21]	C	384 ²	198M	101G	87.5
ConvNeXt-XL [‡] [21]	C	384 ²	350M	179G	87.8
InternImage-L[‡] (ours)	C	384²	223M	108G	87.7
InternImage-XL[‡] (ours)	C	384²	335M	163G	88.0
ViT-G/14 [#] [19]	T	518 ²	1.84B	5160G	90.5
CoAtNet-6 [#] [20]	T	512 ²	1.47B	1521G	90.5
CoAtNet-7 [#] [20]	T	512 ²	2.44B	2586G	90.9
Florence-CoSwin-H [#] [57]	T	—	893M	—	90.0
SwinV2-G [#] [16]	T	640 ²	3.00B	—	90.2
RepLkNet-XL [#] [22]	C	384 ²	335M	129G	87.8
BiT-L-ResNet152x4 [#] [63]	C	480 ²	928M	—	87.5
InternImage-H[#] (ours)	C	224²	1.08B	188G	88.9
InternImage-H[#] (ours)	C	640²	1.08B	1478G	89.6

Table 2. **Image classification performance on the ImageNet validation set.** “type” refers to model type, where “T” and “C” denote transformer and CNN, respectively. “scale” is the input scale. “acc” is the top-1 accuracy. “[‡]” indicates the model is pre-trained on ImageNet-22K [30]. “[#]” indicates pretraining on extra large-scale private dataset such as JFT-300M [64], FLD-900M [57], or the joint public dataset in this work.

and InternImage-L/XL are first trained on ImageNet-22K (~14.2 million) for 90 epochs and then fine-tuned on ImageNet-1K for 20 epochs. To further explore the capability of our model and match the large-scale private data used in previous methods [16, 20, 57], we adopt M3I Pre-training [58], a unified pre-training approach available for both unlabeled and weakly-labeled data, to pre-train InternImage-H on a 427 million joint dataset of public Laion-400M [59], YFCC-15M [60], and CC12M [61] for 30 epochs, and then we fine-tune the model on ImageNet-1K for 20 epochs.

Results. Table 2 shows the classification results of models with different scales. With similar parameters and computational costs, our models are comparable or even su-

prior to the state-of-the-art transformer-based and CNN-based models. For example, InternImage-T achieves 83.5% top-1 accuracy, outperforming ConvNext-T [21] with a clear margin of 1.4 points. InternImage-S/B keeps the leading position and InternImage-B surpasses the hybrid-ViT CoAtNet-2 [20] by 0.8 points. When pre-trained on ImageNet-22K and the large-scale joint dataset, the top-1 accuracy of InternImage-XL and -H are boosted to 88.0% and 89.6%, respectively, which is better than previous CNNs [22, 63] also trained with large-scale data, and closes the gap with the state-of-the-art large-scale ViTs to about 1 point. This gap may be caused by the discrepancy between large-scale inaccessible private data and the aforementioned joint public data. These results show that our InternImage not only has good performance on the common parameter scale and the public training data, but also can effectively extend to large-scale parameters and data.

4.2. Object Detection

Settings. We verify the detection performance of our InternImage on the COCO benchmark [31], on top of two representative object detection frameworks: Mask R-CNN [66], and Cascade Mask R-CNN [67]. We follow common practices [2, 11] to initialize the backbone with pre-trained classification weights, and train models use a $1 \times$ (12 epochs) or $3 \times$ (36 epochs) schedule by default.

Results. As shown in Table 3, when using Mask R-CNN for object detection, we find that under a comparable number of parameters, our models significantly surpass their counterparts. For example, with the $1 \times$ training schedule, the box AP (AP^b) of InternImage-T is 4.5 points better than Swin-T [2] (47.2 vs. 42.7), and 3.0 points higher than ConvNeXt-T [21] (47.2 vs. 44.2). With the $3 \times$ multi-scale training schedule, more parameters, and more advanced Cascade Mask R-CNN [67], InternImage-XL achieves AP^b of 56.2, surpassing ConvNeXt-XL by 1.0 points (56.2 vs. 55.2). Similar results are also seen in instance segmentation experiments. With the $1 \times$ training schedule, InternImage-T yields 42.5 mask AP (*i.e.*, AP^m), which outperforms Swin-T and ConvNeXt-T by 3.2 points (42.5 vs. 39.3) and 2.4 points (42.5 vs. 40.1), respectively. The best AP^m 48.8 is obtained by InternImage-XL with Cascade Mask R-CNN, which is at least 1.1 points higher than its counterparts.

To further push the performance bound of object detection, we follow the advanced setting used in leading methods [16, 17, 26, 70, 74] to initialize the backbone with the weights pre-trained on ImageNet-22K or the large-scale joint dataset, and double its parameters via the composite techniques [74] (see Fig. 2). Then, we fine-tune it along with the DINO [70] detector on the Objects365 [75] and COCO datasets one after another for 26 epochs and 12 epochs, respectively. As shown in Table 4, our method achieves the best results of 65.0 AP^b and 65.4 AP^b on

method	#params	#FLOPs	Mask R-CNN 1× schedule						Mask R-CNN 3×+MS schedule					
			AP ^b	AP ^b ₅₀	AP ^b ₇₅	AP ^m	AP ^m ₅₀	AP ^m ₇₅	AP ^b	AP ^b ₅₀	AP ^b ₇₅	AP ^m	AP ^m ₅₀	AP ^m ₇₅
Swin-T [2]	48M	267G	42.7	65.2	46.8	39.3	62.2	42.2	46.0	68.1	50.3	41.6	65.1	44.9
ConvNeXt-T [21]	48M	262G	44.2	66.6	48.3	40.1	63.3	42.8	46.2	67.9	50.8	41.7	65.0	44.9
PVTv2-B2 [11]	45M	309G	45.3	67.1	49.6	41.2	64.2	44.4	47.8	69.7	52.6	43.1	66.8	46.7
ViT-Adapter-S [65]	48M	403G	44.7	65.8	48.3	39.9	62.5	42.8	48.2	69.7	52.5	42.8	66.4	45.9
InternImage-T (ours)	49M	270G	47.2	69.0	52.1	42.5	66.1	45.8	49.1	70.4	54.1	43.7	67.3	47.3
Swin-S [2]	69M	354G	44.8	66.6	48.9	40.9	63.4	44.2	48.2	69.8	52.8	43.2	67.0	46.1
ConvNeXt-S [21]	70M	348G	45.4	67.9	50.0	41.8	65.2	45.1	47.9	70.0	52.7	42.9	66.9	46.2
PVTv2-B3 [11]	65M	397G	47.0	68.1	51.7	42.5	65.7	45.7	48.4	69.8	53.3	43.2	66.9	46.7
InternImage-S (ours)	69M	340G	47.8	69.8	52.8	43.3	67.1	46.7	49.7	71.1	54.5	44.5	68.5	47.8
Swin-B [2]	107M	496G	46.9	—	—	42.3	—	—	48.6	70.0	53.4	43.3	67.1	46.7
ConvNeXt-B [21]	108M	486G	47.0	69.4	51.7	42.7	66.3	46.0	48.5	70.1	53.3	43.5	67.1	46.7
PVTv2-B5 [11]	102M	557G	47.4	68.6	51.9	42.5	65.7	46.0	48.4	69.2	52.9	42.9	66.6	46.2
ViT-Adapter-B [65]	120M	832G	47.0	68.2	51.4	41.8	65.1	44.9	49.6	70.6	54.0	43.6	67.7	46.9
InternImage-B (ours)	115M	501G	48.8	70.9	54.0	44.0	67.8	47.4	50.3	71.4	55.3	44.8	68.7	48.0

method	#param	#FLOPs	Cascade Mask R-CNN 1× schedule						Cascade Mask R-CNN 3×+MS schedule					
			AP ^b	AP ^b ₅₀	AP ^b ₇₅	AP ^m	AP ^m ₅₀	AP ^m ₇₅	AP ^b	AP ^b ₅₀	AP ^b ₇₅	AP ^m	AP ^m ₅₀	AP ^m ₇₅
Swin-L [‡] [2]	253M	1382G	51.8	71.0	56.2	44.9	68.4	48.9	53.9	72.4	58.8	46.7	70.1	50.8
ConvNeXt-L [‡] [21]	255M	1354G	53.5	72.8	58.3	46.4	70.2	50.2	54.8	73.8	59.8	47.6	71.3	51.7
RepLkNet-31L [‡] [22]	229M	1321G	—	—	—	—	—	—	53.9	72.5	58.6	46.5	70.0	50.6
HorNet-L [‡] [40]	259M	1358G	—	—	—	—	—	—	56.0	—	—	48.6	—	—
InternImage-L [‡] (ours)	277M	1399G	54.9	74.0	59.8	47.7	71.4	52.1	56.1	74.8	60.7	48.5	72.4	53.0
ConvNeXt-XL [‡] [21]	407M	1898G	53.6	72.9	58.5	46.5	70.3	50.5	55.2	74.2	59.9	47.7	71.6	52.2
InternImage-XL [‡] (ours)	387M	1782G	55.3	74.4	60.1	48.1	71.9	52.4	56.2	75.0	61.2	48.8	72.5	53.4

Table 3. **Object detection and instance segmentation performance on COCO val2017.** The FLOPs are measured with 1280×800 inputs. AP^b and AP^m represent box AP and mask AP, respectively. “MS” means multi-scale training.

method	detector	#params	AP ^b		method	crop size	#params	#FLOPs	mIoU (SS)	mIoU (MS)
			val2017	test-dev						
Swin-L [‡] [2]	HTC++ [2]	284M	58.0	58.7	Swin-T [2]	512 ²	60M	945G	44.5	45.8
Swin-L [‡] [2]	Soft-Teacher [68]	284M	60.7	61.3	ConvNeXt-T [21]	512 ²	60M	939G	46.0	46.7
Florence-CoSwin-H [#] [57]	DyHead [69]	637M	62.0	62.4	SLaK-T [29]	512 ²	65M	936G	47.6	—
ViT-Adapter-L [‡] [65]	HTC++ [2]	401M	62.6	62.6	InternImage-T (ours)	512 ²	59M	944G	47.9	48.1
Swin-L [‡] [2]	DINO [70]	218M	63.2	63.3	Swin-S [2]	512 ²	81M	1038G	47.6	49.5
FocalNet-H [‡] [71]	DINO [70]	746M	64.2	64.3	ConvNeXt-S [21]	512 ²	82M	1027G	48.7	49.6
ViT-Huge [72]	Group-DETRv2 [72]	629M	—	64.5	SLaK-S [29]	512 ²	91M	1028G	49.4	—
SwinV2-G [#] [16]	HTC++ [2]	3.00B	62.5	63.1	InternImage-S (ours)	512 ²	80M	1017G	50.1	50.9
BEiT-3 [#] [17]	ViTDet [73]	1.90B	—	63.7	Swin-B [2]	512 ²	121M	1188G	48.1	49.7
FD-SwinV2-G [#] [26]	HTC++ [2]	3.00B	—	64.2	ConvNeXt-B [21]	512 ²	122M	1170G	49.1	49.9
InternImage-XL [‡] (ours)	DINO [70]	602M	64.2	64.3	RepLkNet-31B [22]	512 ²	112M	1170G	49.9	50.6
InternImage-H [#] (ours)	DINO [70]	2.18B	65.0	65.4	SLaK-B [29]	512 ²	135M	1172G	50.2	—
					InternImage-B (ours)	512 ²	128M	1185G	50.8	51.3
					Swin-L [‡] [2]	640 ²	234M	2468G	52.1	53.5
					RepLkNet-31L [‡] [22]	640 ²	207M	2404G	52.4	52.7
					ConvNeXt-L [‡] [21]	640 ²	235M	2458G	53.2	53.7
					ConvNeXt-XL [‡] [21]	640 ²	391M	3335G	53.6	54.0
					InternImage-L [‡] (ours)	640 ²	256M	2526G	53.9	54.1
					InternImage-XL [‡] (ours)	640 ²	368M	3142G	55.0	55.3
					SwinV2-G [#] [16]	896 ²	3.00B	—	—	59.9
					InternImage-H [#] (ours)	896 ²	1.12B	3566G	59.9	60.3
					BEiT-3 [#] [17]	896 ²	1.90B	—	—	62.8
					FD-SwinV2-G [#] [26]	896 ²	3.00B	—	—	61.4
					InternImage-H [#] (ours) + Mask2Former [76]	896 ²	1.31B	4635G	62.5	62.9

Table 4. **Comparison of the state-of-the-art detectors on COCO val2017 and test-dev.**

COCO val2017 and test-dev. Compared to previous state-of-the-art models, we surpass FD-SwinV2-G [26] by 1.2 points (65.4 vs. 64.2), with 27% fewer parameters and without complicated distillation processes, which shows the effectiveness of our models on the detection task.

4.3. Semantic Segmentation

Settings. To evaluate the semantic segmentation performance of InternImage, we initialize the backbone with pre-trained classification weights and train our models with UperNet [77] on ADE20K [78] for 160k iterations and compare fairly with previous CNN-based and transformer-based backbones. To further reach top performance, we arm InternImage-H with more advanced Mask2Former [76], and adopt the same training settings in [17, 65].

Results. As shown in Table 5, when using UperNet [77] for semantic segmentation, our InternImage consistently outperforms prior arts [2, 21, 22, 29]. For exam-

Table 5. **Semantic segmentation performance on the ADE20K validation set.** The FLOPs are measured with 512×2048, 640×2560, or 896×896 inputs according to the crop size.

ple, with almost the same parameter numbers and FLOPs, our InternImage-B reports 50.8 mIoU on the ADE20K val, which is outstanding from the strong counterparts such as ConvNeXt-B (50.8 vs. 49.1) and RepLkNet-31B (50.8 vs. 49.9). Furthermore, our InternImage-H yields 60.3 MS mIoU, which is better than SwinV2-G [16], while the parameter number is much smaller (1.12B vs. 3.00B).

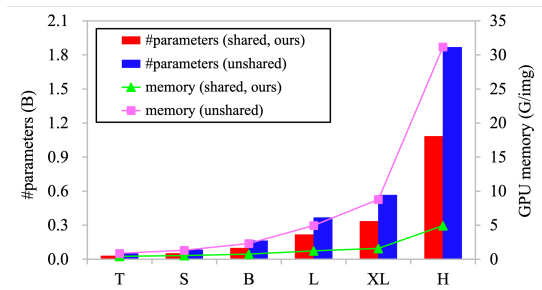


Figure 4. **Model parameters and GPU memory usage of shared weights v.s unshared weights among convolution neurons.** The left vertical axis indicates the model parameters and the right one indicates the GPU memory usage per image when the batch size is 32 and the input image resolution is 224×224 .

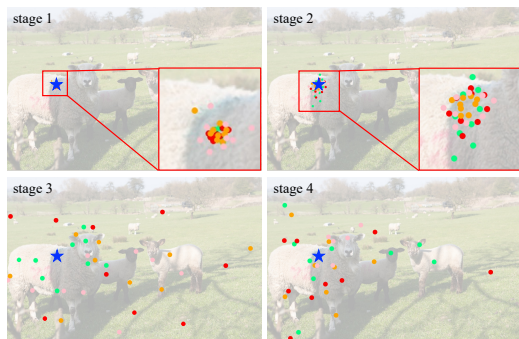


Figure 5. **Visualization of sampling locations for different groups at different stages.** The blue star indicates the query point (on the left sheep), and the dots with different colors indicate the sampling locations of different groups.

It is worth noting that, when using Mask2Former [76] and multi-scale testing, our InternImage-H achieves the best mIoU of 62.9, higher than the current best BEiT-3 [17] on the ADE20K benchmark. These results demonstrate that the CNN-based foundation model can also enjoy the dividends of massive data and challenge the leading position of transformer-based models.

4.4. Ablation Study

Sharing weights among convolution neurons matters.

Large-scale models are sensitive to parameters and memory cost of the core operator, due to hardware limitations. To address this problem, we share weights among convolution neurons of DCNv3. As shown in Fig. 4, we compare the parameters and memory cost of the models based on DCNv3 with shared or unshared weights. We see that the parameters and memory cost of models with unshared weights are much higher than the shared one, especially for the -H scale, the ratio of saved parameters and GPU memory is 42.0% and 84.2%, respectively. As shown in Table 6, we also examine that the two models at -T scale have similar top-1 accuracy on ImageNet (83.5 vs. 83.6) and AP^b on COCO (47.2 vs. 47.4), even the model without shared weights has

shared w	multi-group	softmax	top-1 acc	AP ^b	AP ^m
×	✓	✓	83.6	47.4	42.6
✓	×	✓	82.3	43.8	40.0
✓	✓	×	65.7	38.7	35.6
✓	✓	✓	83.5	47.2	42.5

Table 6. **Ablation comparison of the three modifications in DCNv3.** These experiments are based on InternImage-T for classification and Mask R-CNN $1 \times$ schedule for detection.

66.1% more parameters.

Multi-group spatial aggregation brings stronger features. We introduce aggregation groups to allow our model to learn information from different representation subspaces like transformers [9]. As shown in Fig. 5, for the same query pixel, the offsets from different groups are concentrated in different regions, resulting in hierarchical semantic features. We also compare the performance of the model with and without multiple groups. As reported in Table 6, the model significantly drops 1.2 points on ImageNet and 3.4 points on COCO val2017. In addition, we also see that in the first two stages, the learned effective receptive field (ERF) is relatively small, and as the model goes deeper (*i.e.*, stages 3 and 4), the ERF increases to be global. This phenomenon is different from ViTs [9, 10, 79] whose ERF is usually global. Moreover, the normalization of sampling points improves gradient stability. Without using softmax normalization leads to 17.8 points drop on ImageNet and 8.5 points drop on COCO.

5. Conclusion & Limitations

We introduce InternImage, a new large-scale CNN-based foundation model that can provide strong representations for versatile vision tasks, such as image classification, object detection, and semantic segmentation. We tune the flexible DCNv2 operator to satisfy the requirement of foundation models, and develop a series of blocks, stacking and scaling rules centered on the core operator. Extensive experiments on object detection and semantic segmentation benchmarks verify that our InternImage can obtain comparable or better performance than well-designed large-scale vision transformers trained with massive data, showing that CNN is also a considerable choice for large-scale vision foundation model research. Nonetheless, latency remains an issue for DCN-based operators adapting to downstream tasks with high-speed requirements. Also, large-scale CNNs are still in their early stages of development, and we hope InternImage can serve as a good starting point.

Acknowledgement

This work is partially supported by the National Key RD Program of China (No. 2022ZD0160100), the National Natural Science Foundation of China (Grant No. 61672273, 61832008), and Shanghai Committee of Science and Technology (Grant No. 21DZ1100100).

References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Adv. Neural Inform. Process. Syst.*, 30, 2017. [1](#), [2](#), [3](#), [4](#), [5](#)
- [2] 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 *Int. Conf. Comput. Vis.*, pages 10012–10022, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [3] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019. [1](#)
- [4] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [1](#)
- [5] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020. [1](#)
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Adv. Neural Inform. Process. Syst.*, 33:1877–1901, 2020. [1](#)
- [7] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. [1](#)
- [8] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022. [1](#)
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. Learn. Represent.*, 2020. [1](#), [2](#), [3](#), [5](#), [8](#)
- [10] 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 *Int. Conf. Comput. Vis.*, pages 568–578, 2021. [1](#), [3](#), [5](#), [8](#)
- [11] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):415–424, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [12] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12124–12134, 2022. [1](#)
- [13] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Int. Conf. Comput. Vis.*, pages 22–31, 2021. [1](#)
- [14] Alaaeldin Ali, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, et al. Xcvt: Cross-covariance image transformers. *Adv. Neural Inform. Process. Syst.*, 34, 2021. [1](#)
- [15] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *Adv. Neural Inform. Process. Syst.*, 34, 2021. [1](#)
- [16] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. *Adv. Neural Inform. Process. Syst.*, pages 12009–12019, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [17] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [1](#), [3](#), [6](#), [7](#), [8](#)
- [18] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Adv. Neural Inform. Process. Syst.*, 34:8583–8595, 2021. [1](#), [2](#)
- [19] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12104–12113, 2022. [1](#), [2](#), [3](#), [6](#)
- [20] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *Adv. Neural Inform. Process. Syst.*, 34:3965–3977, 2021. [1](#), [2](#), [3](#), [6](#)
- [21] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *arXiv preprint arXiv:2201.03545*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [22] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11963–11975, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [23] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10819–10829, 2022. [2](#)

- [24] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [2](#), [3](#), [5](#)
- [25] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). arxiv. *arXiv preprint arXiv:1606.08415*, 2016. [2](#), [5](#)
- [26] Yixuan Wei, Han Hu, Zhenda Xie, Zheng Zhang, Yue Cao, Jianmin Bao, Dong Chen, and Baining Guo. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141*, 2022. [2](#), [6](#), [7](#)
- [27] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Int. Conf. Comput. Vis.*, pages 764–773, 2017. [2](#)
- [28] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9308–9316, 2019. [2](#), [3](#), [4](#)
- [29] 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. [2](#), [7](#)
- [30] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009. [2](#), [5](#), [6](#)
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis.*, pages 740–755, 2014. [2](#), [6](#)
- [32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [2](#), [4](#)
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [2](#), [3](#)
- [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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–9, 2015. [2](#)
- [35] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [2](#), [3](#), [5](#)
- [36] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1492–1500, 2017. [2](#)
- [37] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning.*, pages 6105–6114. PMLR, 2019. [2](#), [5](#)
- [38] Mingxing Tan and Quoc Le. Efficientnetv2: Smaller models and faster training. In *International Conference on Machine Learning.*, pages 10096–10106. PMLR, 2021. [2](#)
- [39] 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. [2](#)
- [40] Yongming Rao, Wenliang Zhao, Yansong Tang, Jie Zhou, Ser-Nam Lim, and Jiwen Lu. Hornet: Efficient high-order spatial interactions with recursive gated convolutions. *arXiv preprint arXiv:2207.14284*, 2022. [3](#), [7](#)
- [41] Qi Han, Zejia Fan, Qi Dai, Lei Sun, Ming-Ming Cheng, Jiaying Liu, and Jingdong Wang. On the connection between local attention and dynamic depth-wise convolution. In *Int. Conf. Learn. Represent.*, 2021. [3](#)
- [42] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. [3](#)
- [43] Zhuofan Xia, Xuran Pan, Shiji Song, Li Erran Li, and Gao Huang. Vision transformer with deformable attention. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4794–4803, 2022. [3](#), [4](#)
- [44] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12894–12904, 2021. [3](#)
- [45] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. [3](#)
- [46] Mingyu Ding, Bin Xiao, Noel Codella, Ping Luo, Jingdong Wang, and Lu Yuan. Davit: Dual attention vision transformers. *arXiv preprint arXiv:2204.03645*, 2022. [3](#)
- [47] 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. [3](#)
- [48] Xizhou Zhu, Dazhi Cheng, Zheng Zhang, Stephen Lin, and Jifeng Dai. An empirical study of spatial attention mechanisms in deep networks. In *Int. Conf. Comput. Vis.*, pages 6688–6697, 2019. [3](#)
- [49] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. [3](#)
- [50] L-CCGP Florian and Schroff Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, volume 6, 2017. [3](#)

- [51] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Eur. Conf. Comput. Vis.*, pages 801–818, 2018. 3
- [52] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989. 3
- [53] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1251–1258, 2017. 4
- [54] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 4
- [55] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning.*, pages 10524–10533. PMLR, 2020. 5
- [56] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning.*, pages 10347–10357, 2021. 5
- [57] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 6, 7
- [58] Weijie Su, Xizhou Zhu, Chenxin Tao, Lewei Lu, Bin Li, Gao Huang, Yu Qiao, Xiaogang Wang, Jie Zhou, and Jifeng Dai. Towards all-in-one pre-training via maximizing multi-modal mutual information. In *CVPR*, 2023. 6
- [59] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 6
- [60] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 6
- [61] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3558–3568, 2021. 6
- [62] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022. 6
- [63] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Eur. Conf. Comput. Vis.*, pages 491–507. Springer, 2020. 6
- [64] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10687–10698, 2020. 6
- [65] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022. 7
- [66] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Int. Conf. Comput. Vis.*, pages 2961–2969, 2017. 6
- [67] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: high quality object detection and instance segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(5):1483–1498, 2019. 6
- [68] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *Int. Conf. Comput. Vis.*, pages 3060–3069, 2021. 7
- [69] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7373–7382, 2021. 7
- [70] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 6, 7
- [71] Jianwei Yang, Chunyuan Li, and Jianfeng Gao. Focal modulation networks. *arXiv preprint arXiv:2203.11926*, 2022. 7
- [72] Qiang Chen, Xiaokang Chen, Gang Zeng, and Jingdong Wang. Group detr: Fast training convergence with decoupled one-to-many label assignment. *arXiv preprint arXiv:2207.13085*, 2022. 7
- [73] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. 7
- [74] Tingting Liang, Xiaojie Chu, Yudong Liu, Yongtao Wang, Zhi Tang, Wei Chu, Jingdong Chen, and Haibin Ling. Cbnet: A composite backbone network architecture for object detection. *IEEE Trans. Image Process.*, 2022. 6
- [75] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Int. Conf. Comput. Vis.*, pages 8430–8439, 2019. 6
- [76] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021. 7, 8
- [77] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Eur. Conf. Comput. Vis.*, pages 418–434, 2018. 7

- [78] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 633–641, 2017. [7](#)
- [79] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Adv. Neural Inform. Process. Syst.*, 34, 2021. [8](#)