

Scaling Local Self-Attention for Parameter Efficient Visual Backbones

Ashish Vaswani
Google Research
avaswani@google.com

Prajit Ramachandran
Google Research
prajit@google.com

Aravind Srinivas
UC Berkeley
aravind@cs.berkeley.edu

Niki Parmar
Google Research
parmar@google.com

Blake Hechtman
Google Research
blakehechtman@google.com

Jonathon Shlens
Google Research
shlens@google.com

Abstract

*Self-attention has the promise of improving computer vision systems due to parameter-independent scaling of receptive fields and content-dependent interactions, in contrast to parameter-dependent scaling and content-independent interactions of convolutions. Self-attention models have recently been shown to have encouraging improvements on accuracy-parameter trade-offs compared to baseline convolutional models such as ResNet-50. In this work, we develop self-attention models that can outperform not just the canonical baseline models, but even the high-performing convolutional models. We propose two extensions to self-attention that, in conjunction with a more efficient implementation of self-attention, improve the speed, memory usage, and accuracy of these models. We leverage these improvements to develop a new self-attention model family, HaloNets, which reach state-of-the-art accuracies on the parameter-limited setting of the ImageNet classification benchmark. In preliminary transfer learning experiments, we find that HaloNet models outperform much larger models and have better inference performance. On harder tasks such as object detection and instance segmentation, our simple local self-attention and convolutional hybrids show improvements over very strong baselines. These results mark another step in demonstrating the efficacy of self-attention models on settings traditionally dominated by convolutions.*¹

1. Introduction

Vision and natural language processing (NLP) systems divide the landscape of computational primitives. While self-attention is the primary workhorse in NLP, convolutions

are ubiquitous in nearly all vision models. Convolutions embody the principle of *local* processing, to learn local spatial features such as edges and texture that are abundant in images. On the other hand, the Transformer [53] showed that self-attention is an effective and computationally efficient mechanism for capturing *global* interactions between words in a sentence. Self-attention has several properties that make it a good fit for vision: (a) content-based interactions as opposed to content-independent interactions of convolution; (b) parameter-independent scaling of receptive field size as opposed to parameter-dependent scaling of convolution; (c) empirical ability to capture long-range dependencies for use in larger images; (d) flexibility to handle and integrate multiple types of data that appear in vision, such as pixels [55, 2, 40, 62], point clouds [59], sequence conditioning information [58], and graphs [29]. Self-attention may also be regarded as an adaptive nonlinearity paralleling a long history of techniques in computer vision, such as bilateral filtering [36] and non-local means [3].

Several recent papers [2, 39, 10, 62, 46] have attempted using self-attention primitives to improve image classification accuracy over the strong and commonly used ResNet backbones [14, 15]. Among them, the Stand-Alone Self-Attention (SASA) [39] is a fully self-attentive model that replaces every spatial convolution with *local* self-attention, which improves the performance of ResNet backbones while having fewer parameters and floating point operations. While conceptually promising, these models lag behind state-of-the-art convolutional models in image classification. State-of-the-art convolutional models [51, 63, 38] use a variety of scaling techniques to achieve strong performance across a range of computation and parameter regimes.

In this work, we aim to develop and understand techniques for scaling *local* self-attention models to outperform some of the best convolutional models. Scaling self-attention

¹Please refer to <https://arxiv.org/abs/2103.12731> for a longer version.

models presents a unique set of challenges. For example, convolutions have been very efficiently mapped to matrix accelerators such as TPUs and GPUs that drive most deep learning workloads, but fast implementations of local 2D self-attention do not currently exist. To bridge this gap, we introduce a *non-centered* version of local attention that efficiently maps to existing hardware with *haloing*. While our formulation breaks *translational equivariance*, it improves both throughput and accuracies over the centered local self-attention used in SASA. We also introduce a strided attention downsampling operation for multi-scale feature extraction.

We leverage these techniques to develop a new local self-attention model family, *HaloNet*, which achieves state-of-the-art performance across different parameter regimes. The largest HaloNet achieves 84.9% top-1 accuracy on the ImageNet [43] classification benchmark (Section 4.1). We perform a detailed study to uncover how self-attention and convolutional models scale differently. Our self-attention layers also show promising results on harder tasks such as object detection and instance segmentation (Section 4.5) using the Mask R-CNN framework on the COCO benchmark. Finally, we end with a discussion of current limitations and ideas for future work in applying self-attention to vision.

2. Models and Methods

Although our models use self-attention instead of convolutions for capturing spatial interactions between pixels, they adopt some important architectural features of modern convolutional neural networks (CNNs). Like CNNs, we compute *multi-scale feature hierarchies* [31] which enable detecting objects at multiple sizes in tasks such as localization and instance segmentation. For this, we develop a strided self-attention layer, a natural extension of strided convolutions (Section 2.2). To deal with the computational cost in larger resolutions where global attention is infeasible, we follow the fairly general principle of *local processing*, which is at the heart of convolutions and natural perceptual systems [22, 23], and use spatially restricted forms of self-attention. However, unlike the model of [39], that also use local self-attention, we abstain from enforcing translation equivariance in lieu of better hardware utilization, which improves the speed-accuracy tradeoff (Section 2.2). Also note that while we use local attention, our receptive fields per pixel are quite large (up to 18×18) and we show in Section 4.2.2 that larger receptive fields help with larger images. In the remainder of this section, we will motivate self-attention for vision tasks and describe how we relax translational equivariance to efficiently map local self-attention to hardware.

2.1. Self-attention can generate spatially varying convolutional filters

Self-attention has been viewed as a method to directly capture relationships between distant pixels [39, 19, 54]. It

has also been interpreted as a specific instantiation of the classic technique of non-local means [3, 55]. The perspective that we discuss in this section is one that views self-attention as generating spatially varying filters, in contrast to the reuse of the same filter across every spatial location in standard convolutions [12]. To observe this, we write self-attention and convolution as specific instances of a general spatial pooling function. Given an input $x \in \mathcal{R}^{H \times W \times c_{in}}$, where H is the height, W is the width, and c_{in} is the number of input channels, we define a local 2D pooling function that computes an output at location (i, j) , $y_{ij} \in \mathcal{R}^{c_{out}}$ as

$$y_{ij} = \sum_{a,b \in \mathcal{N}(i,j)} f(i, j, a, b)x_{ab},$$

where $f(i, j, a, b)$ is a function that returns a weight matrix $W \in \mathcal{R}^{c_{in} \times c_{out}}$ at every location in a 2D window $\mathcal{N}(i, j)$ of size $k \times k$ centered at (i, j) . Note that later in this section, we introduce non-centered windows for self-attention, but we use centering here for ease of explanation. This computation is repeated for every pixel (i, j) . For a convolution, $f(i, j, a, b)$ returns a *different linear transformation* for each relative distance in neighborhood, and these weights are shared across all (i, j) . Weight sharing significantly reduces parameters and encourages learning features that repeat spatially. In dot-product relative self-attention [44, 39, 2] (eqs. (2) and (3)), every pixel in the neighborhood shares the *same linear transformation* which is multiplied by a scalar probability that is a function of both content-content and content-geometry interactions resulting in weights that can vary spatially. As an example, for a ball and an orange at two different locations in an image, pixels inside the ball and the orange are likely to generate different $p_{a-i, b-j}^{ij}$ because of the different content around them, such as color or texture.

$$f(i, j, a, b)^{conv} = W_{a-i, b-j} \quad (1)$$

$$f(i, j, a, b)^{self-att} = \text{softmax}_{ab} \left((W_Q x_{ij})^\top W_K x_{ab} + (W_Q x_{ij})^\top r_{a-i, b-j} \right) W_V \quad (2)$$

$$= p_{a-i, b-j}^{ij} W_v \quad (3)$$

For self-attention, W_Q , W_K , and W_V are learned linear transformations that are shared across all spatial locations, and respectively produce *queries*, *keys*, and *values* when used to transform x . Spatial geometry is captured by $r_{a-i, b-j}$, which is a learned relative position based embedding. The $(W_Q x_{ij})^\top W_K x_{ab}$ component captures the content-content interaction between the query pixel and a key pixel in the window. The $(W_Q x_{ij})^\top r_{a-i, b-j}$ component is the content-geometry interaction that captures the relationship between the query and the relative position of the key pixel [44]. Note

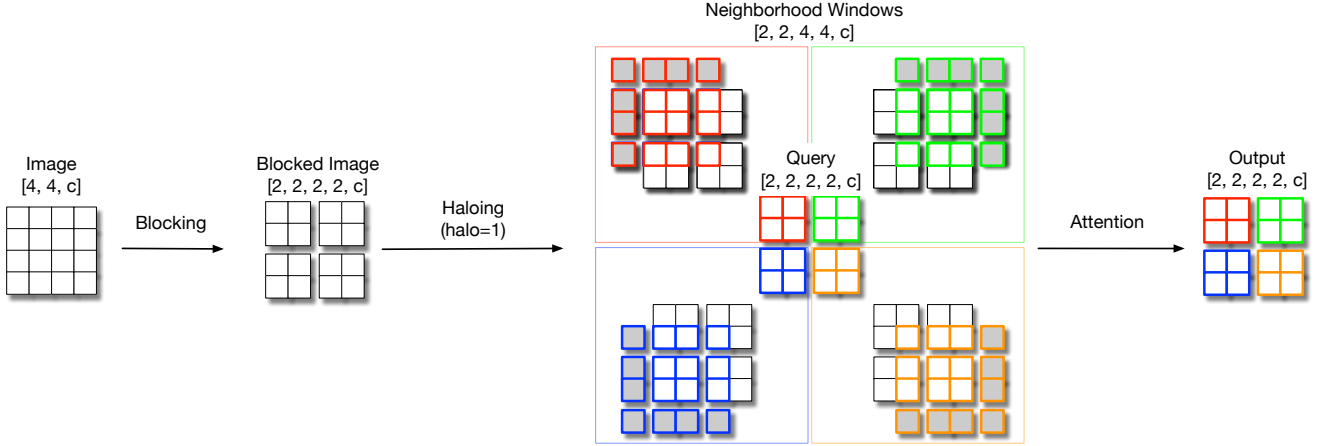


Figure 1. **HaloNet local self-attention architecture:** The different stages of blocked local attention for a $[4, 4, c]$ image, block size $b = 2$, and halo $h = 1$. The image is first blocked into non-overlapping $[2, 2, c]$ images from which the queries are computed. The subsequent haloing step then extracts a $[4, 4, c]$ memory around each of the blocks which linearly transform to keys and values. The spatial dimensions after attention are the same as the queries.

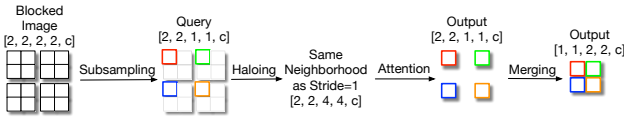


Figure 2. The attention downsampling layer subsamples the queries but keeps the neighborhood the same as the the stride=1 case.

that this formulation preserves *translational equivariance*. If an object translates in an image, for any pixel within the object, the content around it stays the same, generating the same $p_{a-i, b-j}^{ij}$, thereby producing the same output after self-attention. To increase expressivity, multi-headed attention [53] is used, which repeats this computation multiple times in parallel with different parameters, analogous to group convolutions [27, 57].

In the SASA model of [39], the local window $\mathcal{N}(i, j)$ is a $k \times k$ window centered around (i, j) , just like a convolution. The size of this local window k is an important setting to leverage in self-attention. Unlike dense convolutions, k can grow without significantly increasing the number of parameters. Since the projection parameters (W_Q, W_K, W_V) are independent of k , the only parameters that increase with k is $r_{a-i, b-j}$. However, $r_{a-i, b-j}$ constitutes a trivial fraction of the parameters compared to the projection parameters², so increasing k does not impact the number of parameters of the layer significantly. In contrast, the number of parameters in a convolution layer scale quadratically with k (e.g., a 5×5 convolution has $\frac{25}{9}$ times the parameters of a 3×3

²For a window size as large as 63, and 16 dimensions per attention head, $r_{a-i, b-j}$ would add only $63 * 16 = 1008$ parameters per layer because $r_{a-i, b-j}$ are shared among heads. In contrast, if the dimensions of the attention layer were 512, W_Q, W_K, W_V would contribute 786432 parameters. We show details in the appendix.

convolution). On the other hand, the computational cost of self-attention grows quadratically with k , preventing the use of very large values for k .

2.2. Improving the speed-memory tradeoff by relaxing translational equivariance

Global self-attention, in which all locations attend to each other, is too expensive for most image scales due to the quadratic computation cost with respect to k . Thus, multi-scale visual backbones need to use local attention to limit the size of k . We follow the intuitive form of local attention developed in [39], which tries to mimic the square neighborhoods used by convolutions. This form of local attention requires extracting local 2D grids around each pixel. Unfortunately, while deep learning libraries automatically handle neighborhood gathering for convolutions, no such neighborhood gathering function exists for local self-attention (or any general local function). Thus, implementing local self-attention requires explicitly gathering the local neighborhoods before the actual self-attention operation can be performed. While the implementation of this local neighborhood gathering function might initially appear to be a relatively minor implementation detail, in practice, it must actually be carefully designed to reduce memory usage while avoiding unnecessary extra computation. An unoptimized implementation can prevent self-attention models from scaling up due to either out-of-memory errors or excessive slowness. The following discussion frames the design considerations of this neighborhood gathering function.

A straightforward approach would gather $k \times k$ sized windows *separately* around each pixel. As summarized in Table 1 (Row 1), this method blows up the memory used by a factor of k^2 due to replicating the pixel contents for each of

Method	Neighborhood Memory	Receptive Field	FLOPs Per Pixel
Global	HWc	$HW \times HW$	$4(HW)^2c$
Per pixel windows	HWk^2c	$k \times k$	$4k^2c$
SASA [39]	$\frac{HW}{b^2}(b+2h)^2c$	$k \times k$, where $h = \lfloor \frac{k}{2} \rfloor$	$4(b+2h)^2c$
Blocked local (ours)	$\frac{HW}{b^2}(b+2h)^2c$	$(b+2h) \times (b+2h)$	$4(b+2h)^2c$

Table 1. **Scaling behavior of self-attention mechanisms.** f is the number of heads, b is the size of the block, c is the total number of channels, and h is the size of the halo

the k^2 neighborhoods it participates in. This solution quickly leads to out-of-memory errors. *Global* attention (Row 4) is at the other end of the spectrum, where all pixels *share* the same neighborhood, lowering memory at the expense of considerably more FLOPs³. This solution slows down models significantly, while also imposing memory problems due the massize size of the attention matrix. A solution that lies in-between these two extremes should trade-off memory and compute appropriately, with the recognition that a small amount of waste is required.

A compromise solution can be achieved by leveraging the idea that *neighboring pixels share most of their neighborhood*. For example, two pixels that are right next to each other share $k \times (k - 1)$ pixels of their neighborhoods. Thus a local neighborhood for a *block* of pixels can be extracted once together, instead of extracting separate neighborhoods per pixel. The FLOPs can be controlled by varying the number of pixels that form a block. We name this strategy *blocked local self-attention*. The two extremes discussed above are a special case of blocked local self-attention. Global attention corresponds to setting the block size to be the entire spatial extent, while the per-pixel extraction corresponds to setting the block size to be 1.

Figure 1 depicts the different steps involved in executing blocked local self-attention for an image with height $H = 4$, width $W = 4$, and c channels with stride 1. Blocking chops up the image into a $\frac{H}{b}, \frac{W}{b}$ tensor of *non-overlapping* (b, b) blocks. Each block behaves as a group of query pixels and a *haloing* operation combines a band of h pixels around them (with padding at boundaries) to obtain the corresponding *shared neighborhood* block of shape $(\frac{H}{b}, \frac{W}{b}, b + 2h, b + 2h, c)$ from which the keys and values are computed. $\frac{H}{b} \times \frac{W}{b}$ attention operations then run in parallel for each of the query blocks and their corresponding neighborhoods, illustrated with different colors in Figure 1. SASA [39] used the same blocking strategy⁴, setting $h = \lfloor \frac{k}{2} \rfloor$ and uses attention masks to emulate pixel-centered neighborhood windows of size $k \times k$. Our approach For example, to achieve a 7×7 pixel

³To illustrate this, on a 128×128 resolution with 64 channels, global self-attention would incur about 28 times more FLOPs than a 3×3 convolution with 64 input and output channels

⁴Code for both SASA and HaloNet will be made available, along with the checkpoints for HaloNet

centered window, [39] set $h = 3$. The use of attention masks gives the operation translational equivariance, since each pixel only looks at a square window around it.

However, the downside of using attention masks is that it wastes computation that must happen regardless due to the implementation of this algorithm. If attention masks are not used, the receptive field increases without any additional computation, as shown in Table 1 (Rows 2 and 3). However, pixel-level translational equivariance is lost because the non-square receptive fields means that the output of a pixel is dependent on which block it falls into. Take for example a pixel at the left edge of its block, which sees additional pixels that are to the right of its square receptive field. If the entire image is shifted one pixel to the right, the pixel now falls into right edge of a neighboring block, and now sees additional pixels that are to the left of its square receptive field. Thus the output of the pixel is dependent on its position in a block, which can change if the image shifts. Another perspective is that blocked local self-attention is only translational equivariant to shifts of size b . While pixel-level translational equivariance is considered important for achieving good performance[61], we find that empirically, using a non-masked block local self-attention actually improves the accuracy of the model (see Section 4.3). We suspect that the image shifting and cropping perturbations in common data augmentation strategies reduce the reliance on such inductive biases. Thus we adopt unmasked blocked local self-attention because it improves accuracy without sacrificing performance.

Another difference with SASA is our implementation of downsampling. We replace attention followed by post-attention strided average pooling by a single strided attention layer that subsamples queries similar to strided convolutions, as shown in Figure 2. Note that we use the same neighborhood as is extracted in the stride 1 case (Figure 1). This change does not impact accuracy while also reducing the FLOPs $4 \times$ in the downsampling layers. We also implement some important algorithmic optimizations that improve our throughput primarily by avoiding reshapes and data formatting operations. In interest of space, we list them in the Appendix D. Taken together, the speedups produced by these improvements are significant as seen in Figure 3, with up to $2 \times$ improvements in step time. These improvements can be leveraged to train large self-attention models that were previously too expensive. We leave additional optimizations, such as fused operations and better pipelining of memory accesses with computation, to future work.

Note that in the deeper layers of multiscale architectures, smaller spatial dimensions and larger channels would shift the compute calculus in favor of global attention. The models we introduce in Section 4, also take advantage of this, typically using local attention in the higher resolutions and global attention when the image resolutions are the smallest.

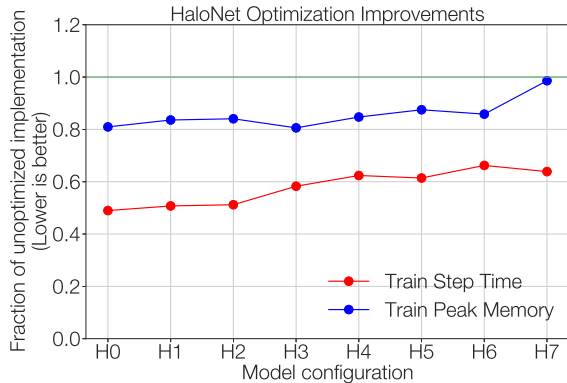


Figure 3. **Optimizations improve performance.** The improvements here are a result of reducing FLOPs with our attention downsampling and improved local self-attention algorithms that avoid reshapes and data formatting. In some cases, we halve the training step time computed on TPU v3.

Output Resolution	Layers
$\frac{s}{4} \times \frac{s}{4}$	7×7 conv stride 2, 64 3×3 max pool stride 2
$\frac{s}{4} \times \frac{s}{4}$	$\left\{ \begin{array}{l} 1 \times 1, 64 \\ \text{attention}(b, h), 64 \cdot r_v \\ 1 \times 1, 64 \cdot r_b \end{array} \right\} \times 3$
$\frac{s}{8} \times \frac{s}{8}$	$\left\{ \begin{array}{l} 1 \times 1, 128 \\ \text{attention}(b, h), 128 \cdot r_v \\ 1 \times 1, 128 \cdot r_b \end{array} \right\} \times 3$
$\frac{s}{16} \times \frac{s}{16}$	$\left\{ \begin{array}{l} 1 \times 1, 256 \\ \text{attention}(b, h), 256 \cdot r_v \\ 1 \times 1, 256 \cdot r_b \end{array} \right\} \times l_3$
$\frac{s}{32} \times \frac{s}{32}$	$\left\{ \begin{array}{l} 1 \times 1, 512 \\ \text{attention}(b, h), 512 \cdot r_v \\ 1 \times 1, 512 \cdot r_b \end{array} \right\} \times 3$
$\frac{s}{32} \times \frac{s}{32}$	$1 \times 1, d_f$
1×1	global average pooling fc, 1000

Table 2. **HaloNet model family specification.**

2.3. HaloNet

Using the implementation of local 2D self-attention with haloing detailed above, we propose a new model, *HaloNet* that matches state-of-the-art convolutional models on the parameter-accuracy trade-off curve. We leverage the structure of ResNets [14] that stack multiple residual bottleneck blocks together (see Table 2). HaloNet uses a few minor modifications from ResNets: (a) adding a final 1×1 convolution before the global average pooling for larger models, following EfficientNet [51], (b) modifying the bottleneck block width factor, which is traditionally fixed at 4, (c) mod-

ifying the output width multiplier of the spatial operation, which is traditionally fixed at 1, (d) changing the number of blocks in the third stage from 4 to 3 for computational reasons because attention is more expensive in the higher resolution layers. We also fix the number of heads for each of the four stages to (4, 8, 8, 8) because heads are more expensive at higher resolutions. To summarize, the scaling dimensions in HaloNet are: image size s , query block size b , halo size h , attention output width multiplier r_v , bottleneck output width multiplier r_b , number of bottleneck blocks in the third group l_3 , and final 1×1 conv width d_f . Our attention neighborhoods range from 14×14 ($b = 8, h = 3$) to 18×18 ($b = 14, h = 2$).

Since the ResNet structure was initially designed for convolutions, we suspect that designing architectures specifically for attention may improve HaloNet. In our work we maintained homogeneity across all layers of model for hyperparameters such as the block (b) and halo (h) sizes. We also hope that using automated architecture search methods [51] to optimize these hyperparameters for specific accelerators will lead to better local attention architectures. In our work, we train with comparable image sizes as EfficientNet models to determine if attention models can scale to larger images.

3. Related Work

We directly build on top of the approach of [39], who compute attention on local regions in order to build a fully self-attentional vision model for classification and object detection. Different forms of attention for pure self-attention vision models have also been proposed [19, 62], which are orthogonal and complementary to the focus on scaling in this work. In addition to attention over the spatial extent that we focus on, components that perform attention over channels have also been used to augment convolutional models [20, 30]. In recent and concurrent work, Vision Transformer [10] show that applying transformers on projections of *non-overlapping* image patches can achieve accuracies comparable to SOTA when pre-trained on very large (JFT-300M [48]) and medium sized (ImageNet-21k [9]) classification datasets. However, their models do not adopt a multiscale architecture and our focus in this work is training on ImageNet [43] from scratch. In Section C.3, we conduct transfer experiments and compare with ViT and BiT [26].

Generally, the performance of computational primitives tend to improve over time due to algorithmic changes to the primitive and better software implementations. Convolutions have improved over the last decade through changes in (a) the computation of the primitive [4, 25, 34, 52, 56, 28]; (b) the software implementation [5]; (c) the structure of the primitive itself, through for example, grouped convolution [57] and depthwise separable convolution [45]. Attention is in the beginning phases of this performance improvement trajectory, and given its importance in sequence modeling [53], it

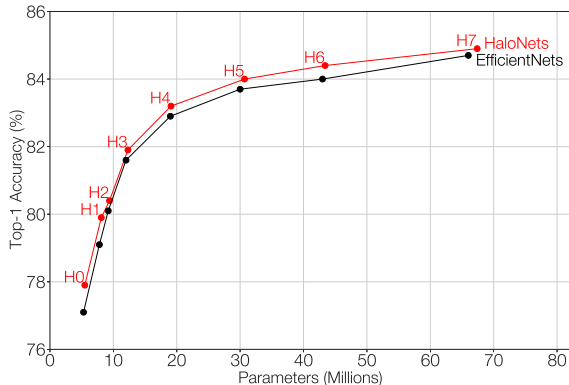


Figure 4. **HaloNets can match EfficientNets on the accuracy vs. parameter trade-off.** The accuracies for EfficientNets B5 and B7 were obtained using RandAugment. For a discussion on training speed, please see Section E.

will likely see sustained effort to enhance performance. Local attention could also receive performance improvements if it is adopted more widely to combat the general problem of processing large inputs. Our work introduces blocked local attention to efficiently process immediate neighbors. Other forms of non-global pixel interaction can also be implemented efficiently [6, 18, 54, 1].

4. Experiments

Each HaloNet model (H0–H7) is designed by successively growing the values of the hyperparameters defined in Table 2. In interest of space, we leave the exact configurations of our models to the Appendix C.1. We also leave the training and evaluation of larger HaloNet models that compare with larger EfficientNet models for future work.

4.1. HaloNets are competitive with state-of-the-art convolutional models

We train our HaloNet models on ImageNet [43] (ILSVRC-2012) benchmark with a batch size of 4096 and learning rate of 1.6, which is linearly warmed up for 10 epochs and followed by cosine decay [33]. The models are trained for 350 epochs with Nesterov’s Accelerated Gradient [35, 49], and regularized with dropout [47], weight decay, RandAugment [8] and stochastic depth [21].

We find that HaloNets perform at par or slightly better (Figure 4) than EfficientNet models for the same parameters, outperforming other model families. Our best model, H7, achieves **84.9%** top-1 ImageNet validation accuracy and **74.7%** top-1 accuracy on ImageNet V2 [41] (with a -0.5% gap to the linear fit in [41]). For each of our HaloNet models, we use image sizes comparable to the corresponding EfficientNet model, training on images sizes up to 600×600 . (Table A2). For a comparison of our latencies with EfficientNet, the reader can refer to Section 5. To the best of

our knowledge, these results are the first to show that self-attention based models for vision perform on par with the SOTA for image classification when trained on imagenet from scratch. Note that for all our experiments, we report accuracies at the end of training and we tune regularization hyperparameters such as augmentation hyperparameters for the baselines and HaloNet models.

4.2. Model study 1: comparing self-attention and convolutions

In the following sections, we will focus on model studies to distinguish the advantages of self-attention over convolutions for vision and understand how to best design self-attention vision architectures. This knowledge is important since much of the progress in convolutional networks comes from improvements in architecture design while keeping the core convolution primitive the same [27, 50, 14]. We believe our study is the first to explicitly examine the design of optimal self-attention vision architectures.

For the remainder of the experimental section, we compare with ResNet-50 [15], the canonical vision model, because many of the components that we ablate have been well studied for ResNet-50, allowing us to use best practices for the baseline model. We tune our baseline ResNet-50 implementation to achieve a better accuracy, 77.6%, compared to commonly reported numbers in the literature. For example, [14] report 76.3%. We then create a new HaloNet architecture, HaloNet-50, that exactly matches the ResNet-50 architecture by replacing spatial convolutions with local self-attention. HaloNet-50 and ResNet-50 have about 18.0 million and 25.5 million parameters respectively. We train both for 150 epochs on 256×256 image size. We share other training details of the ablation set-up in the appendix

4.2.1 Transfer of convolutional components to self-attention

Utilizing regularizations and architectural modules beyond the core primitive is critical for achieving strong results [16]. In this section, we study the effects of these additional components on self-attention models. The components we study were all designed for use in convolutional models, as they were developed through experimentation (either human or automated search) on convolutional models. We examine whether these components can successfully transfer to the new model family of self-attention networks.

We focus on 4 different components based on the design of EfficientNet [51], 2 architecture modules and 2 regularizations: Squeeze-and-Excitation (SE) [20], a channel attention module used after the spatial convolution; SiLU/Swish-1 [40, 11, 17], an activation function with the form $x \cdot \text{sigmoid}(x)$; RandAugment (RA) [8], a data augmentation scheme that simplifies AutoAugment [7]; and Label

Smoothing (LS) [50], a smoothing of the label distribution.

The results from adding these components to the baseline model are in Table 3. Surprisingly, regularizations of the same strength improve HaloNet accuracies significantly more than ResNet, despite HaloNet having around 30% fewer parameters than ResNet. When label smoothing and RandAugment are added, HaloNet improves by 1.3% while ResNet improves by 0.8%. This suggests that self-attention models may require regularizations that are typical of larger convolutional models, perhaps due to the expressivity of self-attention.

When Squeeze-and-Excitation (SE) and SiLU/Swish-1 are added, ResNet improves by 1.3% while HaloNet only improves by 0.4%. We speculate that HaloNet models benefit from the gating and multiplicative interactions that comprise self-attention and do not need explicit gating such as SE. Further research must be conducted in order to discover architecture modules that can consistently improve a variety of self-attention models. Inspired by these findings, we decided to use label smoothing, SiLU/Swish-1, and RandAugment in our HaloNet $H_0 - H_7$ models. We also use stochastic depth for our larger models [21, 51].

Components	HaloNet Accuracy	Baseline Δ	ResNet Accuracy	Baseline Δ
Baseline	78.6	0.0	77.6	0.0
+ LS	79.7	1.1	78.1	0.5
+ LS, RA	79.9	1.3	78.4	0.8
+ SE	78.6	0.0	78.6	1.0
+ SE, SiLU/Sw1	79.0	0.4	78.9	1.3
+ LS, SE	79.7	1.1	78.9	1.3
+ LS, SE, SiLU/Sw1	79.9	1.3	79.1	1.5
+ LS, SE, SiLU/Sw1, RA	80.5	1.9	79.5	1.9

Table 3. **HaloNet improves more than ResNet with regularizations, but does not improve significantly with architectural modules that strongly benefit ResNet.** Starting from a baseline model, adding label smoothing (LS), RandAugment (RA), Squeeze-and-Excitation (SE), and SiLU/Swish-1 (SiLU/Sw1).

4.2.2 Increasing image sizes improve accuracies

A beneficial property of self-attention is attention is that the receptive field size can scale along with image size without significantly impacting the number of parameters (see Section 2.1). As shown in Figure 6, HaloNet consistently improves when using larger images. Although we also see improvements with convolutional models, the accuracy gap between HaloNets and ResNets is maintained.

4.3. Model study 2: HaloNet architecture study

In this section, we will study the impact of relaxing translational equivariance and the relationship of neighborhood window and halo sizes. In the interest of space, a detailed

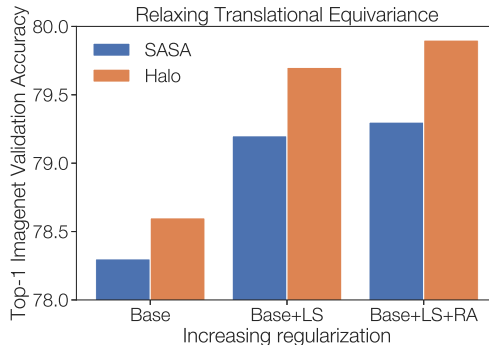


Figure 5. **Relaxing translational equivariance improves accuracies**

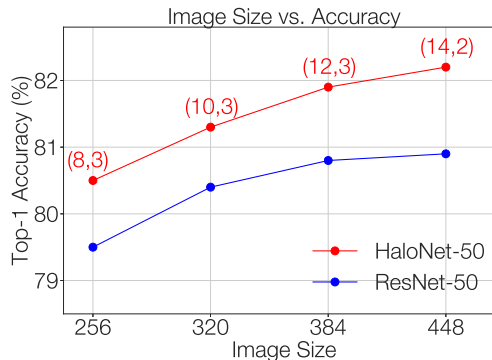


Figure 6. **The accuracy gap between HaloNet-50 and ResNet-50 is maintained with increasing image sizes.** The HaloNet experiments are annotated with block size (b), halo size (h).

study of scaling various components of our models such as r_v , r_{qk} etc can be found in the Appendix B.

Relaxing translational equivariance: In Figure 5, we see that HaloNet-50 with $b = 8$, and $h = 3$ achieves better accuracies using the same block and halo to achieve 7×7 neighborhoods with attention masks [39] and the gap widens with more regularizations. This suggests that larger receptive fields are more important than inductive biases such as translational equivariance.

Window and halo size: When using the blocked input format, there are two ways of changing the window size of attention: changing the query block size or the halo size. For the same window size w , smaller query blocks and larger halos require more memory than larger query blocks and smaller halos, as discussed in section 2.2.

We see in Figure 7 that accuracy consistently improves as the window size increases. In particular, doubling the window size from 6×6 to 12×12 produces a 1.3% accuracy gain. These results suggest that increasing window size can be successfully used to scale models without increasing the

Conv Stages	Attention Stages	Top-1 Acc (%)	Norm. Train Time
-	1, 2, 3, 4	84.9	1.9
1	2, 3, 4	84.6	1.4
1, 2	3, 4	84.7	1.0
1, 2, 3	4	83.8	0.5

Table 4. **Replacing attention layers with convolutions in stages 1 and 2 exhibit the best speed vs. accuracy tradeoff.** All the models had about 67 million parameters and the train and inference times are normalized to the corresponding times for EfficientNet B7. Please see Figure A1 for a detailed comparison of step times.

number of parameters, potentially beneficial for production environments. Furthermore, for a fixed window size, the choice of query block size does not impact results, enabling the usage of larger query block sizes to reduce memory. Figure 7 also shows that eschewing haloing for *non-overlapping* attention, can lower accuracy significantly unless the blocks are quite large. For example using a block size of 4 and a halo of 1 results in better accuracy than using a block size of 8 with 0 halo, despite a smaller neighborhood size.

4.4. Convolution-Attention hybrids improve the speed-accuracy tradeoff

In our final set of ablations, we replace self-attention with convolutions to understand where attention layers are currently most beneficial. In Table 4, we show results for replacing attention layers with convolutions with squeeze-and-excitation modules in each of the stages of our best performing model (HaloNet H7). Having convolutions in all stages except the last yields the fastest model albeit with a significant loss in top-1 accuracy (1%). Splitting the allocation between convolutions (in stages 1–2) and attention (in stages 3–4) minimally detracts predictive accuracy while significantly improving training and inference step times. We leave a detailed study of improved hybrid models for future work.

Model	AP ^{bb}	AP ^{mk}	Speed (ms)
R50 baseline in lit	42.1	37.7	409
R50 + SE (our baseline)	44.5 (+2.4)	39.6 (+1.9)	446
R50 + SE + Local Att ($b = 8$)	45.2 (++)	40.3 (++)	540
R50 + SE + Local Att ($b = 32$)	45.4 (++)	40.5 (++)	613
R101 + SE (our baseline)	45.9 (+3.8)	40.6 (+2.9)	740
R101 + SE + Local Att ($b = 8$)	46.8 (++)	41.2 (++)	799

Table 5. **Accuracies on object detection and instance segmentation.** *bb* (bounding box) refers to detection, and *mk* (mask) refers to segmentation. Speed is measured as the milliseconds taken by only the backbone (and not the FPN) for a batch size of 32 on 2 TPUv3 cores. Please find detailed accuracies in Table A4.

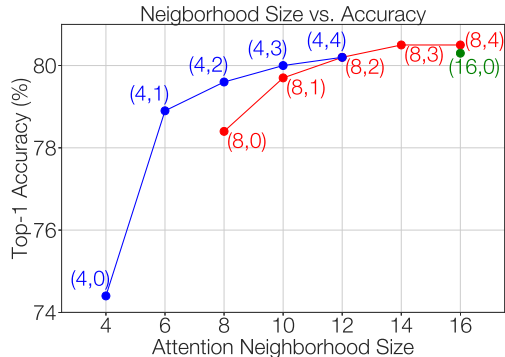


Figure 7. **Increasing window sizes improves accuracy up to a point.** The experiments in the graph have been annotated with their block size (b), halo size (h), $h = 0$ implies attention with *non-overlapping* blocks

4.5. Detection and instance segmentation

To understand if our primitives will generalize to structured prediction tasks on larger images, we conduct initial investigations with the simple attention-convolutional hybrids on detection and instance segmentation, using the Mask R-CNN [13] framework. These hybrids are also faster and consume less memory than pure attention models, enabling faster experimental cycles. We provide more training details in the Appendix C.4.

Our ResNet-50 baseline in row 2 of Table 5, is significantly better than what is usually reported in the literature (row 1). Our attention variants achieve at least 0.7 mAP gains on bounding box detection and at least 0.6 mAP gains on instance segmentation on top of our stronger baselines (denoted by ++ in rows 3, 4 and 6 in Table 5). The gain from local attention with block size $b = 8$ closes half of the mAP gap between the R50 and R101 baselines in detection and 70% of the gap in instance segmentation despite being less than a third of the gap in terms of wall-clock time. Local attention with $b = 8$ and $h = 3$ also improves on top of the deep R101 backbone. These models have only three layers of self-attention, and more layers could alter these results. We leave the study of detection and instance segmentation with pure attention models to future work.

5. Conclusion

In this work, we built multiscale self-attention models that are competitive with the best convolutional models. To achieve this result, we developed two attention improvements: blocked local attention and attention downsampling. Overall, our work shows that self-attention can be competitive in regimes traditionally dominated by computer vision. Future work can push these boundaries further, both in terms of scale and efficiency.

References

- [1] Irwan Bello. Lambdanetworks: Modeling long-range interactions without attention. *arXiv preprint arXiv:2102.08602*, 2021.
- [2] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3286–3295, 2019.
- [3] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 60–65. IEEE, 2005.
- [4] Kumar Chellapilla, Sidd Puri, and Patrice Simard. High performance convolutional neural networks for document processing. In *Tenth International Workshop on Frontiers in Handwriting Recognition*, 2006.
- [5] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759*, 2014.
- [6] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [7] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 113–123, 2019.
- [8] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical data augmentation with no separate search. *arXiv preprint arXiv:1909.13719*, 2019.
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [10] 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. *arXiv preprint arXiv:2010.11929*, 2020.
- [11] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.
- [12] Gamaleldin F Elsayed, Prajit Ramachandran, Jonathon Shlens, and Simon Kornblith. Revisiting spatial invariance with low-rank local connectivity. *arXiv preprint arXiv:2002.02959*, 2020.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [16] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019.
- [17] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [18] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
- [19] Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3464–3473, 2019.
- [20] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [21] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016.
- [22] David H Hubel and TN Wiesel. Shape and arrangement of columns in cat’s striate cortex. *The Journal of physiology*, 165(3):559, 1963.
- [23] David H Hubel and Torsten N Wiesel. Receptive fields and functional architecture of monkey striate cortex. *The Journal of physiology*, 195(1):215–243, 1968.
- [24] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 675–678, 2014.
- [26] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. *arXiv preprint arXiv:1912.11370*, 6(2):8, 2019.
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [28] Andrew Lavin and Scott Gray. Fast algorithms for convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4013–4021, 2016.
- [29] Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. Relation-aware graph attention network for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10313–10322, 2019.
- [30] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 510–519, 2019.

- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [32] 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 *European conference on computer vision*, pages 740–755. Springer, 2014.
- [33] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [34] Michael Mathieu, Mikael Henaff, and Yann LeCun. Fast training of convolutional networks through ffts. *arXiv preprint arXiv:1312.5851*, 2013.
- [35] Y. E. NESTEROV. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. *Dokl. Akad. Nauk SSSR*, 269:543–547, 1983.
- [36] Sylvain Paris, Pierre Kornprobst, Jack Tumblin, and Frédo Durand. *Bilateral filtering: Theory and applications*. Now Publishers Inc, 2009.
- [37] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- [38] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. *arXiv preprint arXiv:2003.13678*, 2020.
- [39] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 68–80. Curran Associates, Inc., 2019.
- [40] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [41] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019.
- [42] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *arXiv preprint arXiv:2003.05997*, 2020.
- [43] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [44] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*, 2018.
- [45] Laurent Sifre. Rigid-motion scattering for image classification. *Ph. D. thesis*, 2014.
- [46] Aravind Srinivas, Tsung-Yi Lin, Niki Parmar, Jonathon Shlens, Pieter Abbeel, and Ashish Vaswani. Bottleneck transformers for visual recognition. *arXiv preprint arXiv:2101.11605*, 2021.
- [47] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [48] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017.
- [49] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*, 2013.
- [50] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [51] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [52] Nicolas Vasilache, Jeff Johnson, Michael Mathieu, Soumith Chintala, Serkan Piantino, and Yann LeCun. Fast convolutional nets with fbfft: A gpu performance evaluation. *arXiv preprint arXiv:1412.7580*, 2014.
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [54] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. *arXiv preprint arXiv:2003.07853*, 2020.
- [55] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [56] Shmuel Winograd. *Arithmetic complexity of computations*, volume 33. Siam, 1980.
- [57] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [58] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015.
- [59] Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, and Qi Tian. Modeling point clouds with self-attention and gumbel subset sampling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3323–3332, 2019.
- [60] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xi-aodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. *arXiv preprint arXiv:2003.11142*, 2020.

- [61] Richard Zhang. Making convolutional networks shift-invariant again. *arXiv preprint arXiv:1904.11486*, 2019.
- [62] Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun. Exploring self-attention for image recognition. *arXiv preprint arXiv:2004.13621*, 2020.
- [63] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018.