

Bootstrapping SparseFormers from Vision Foundation Models

Ziteng Gao

Zhan Tong

Kevin Qinghong Lin

Joya Chen

Mike Zheng Shou✉

Show Lab, National University of Singapore

Abstract

The recently proposed SparseFormer architecture provides an alternative approach to visual understanding by utilizing a significantly lower number of visual tokens via adjusting RoIs, greatly reducing computational costs while still achieving promising performance. However, training SparseFormers from scratch is still expensive, and scaling up the number of parameters can be challenging. In this paper, we propose to bootstrap SparseFormers from ViT-based vision foundation models in a simple and efficient way. Since the majority of SparseFormer blocks are the standard transformer ones, we can inherit weights from large-scale pre-trained vision transformers and freeze them as much as possible. Therefore, we only need to train the SparseFormer-specific lightweight focusing transformer to adjust token RoIs and fine-tune a few early pre-trained blocks to align the final token representation. In such a way, we can bootstrap SparseFormer architectures from various large-scale pre-trained models (e.g., IN-21K pre-trained AugRegs or CLIPs) using a rather smaller amount of training samples (e.g., IN-1K) and without labels or captions within just a few hours. As a result, the bootstrapped unimodal SparseFormer (from AugReg-ViT-L/16-384) can reach 84.9% accuracy on IN-1K with only 49 tokens, and the multimodal SparseFormer from CLIPs also demonstrates notable zero-shot performance with highly reduced computational cost without seeing any caption during the bootstrapping procedure. In addition, CLIP-bootstrapped SparseFormers, which align the output space with language without seeing a word, can serve as efficient vision encoders in multimodal large language models. Code and models are available at <https://github.com/showlab/sparseformer>

1. Introduction

Large-scale pre-trained vision models [13, 28, 49, 59, 76], or vision foundation models, exhibit strong transferable or zero-shot capabilities after pre-training. Most vision foundation models, e.g., ViTs [16], are simply based on the standard transformer encoder architecture [15], which is known

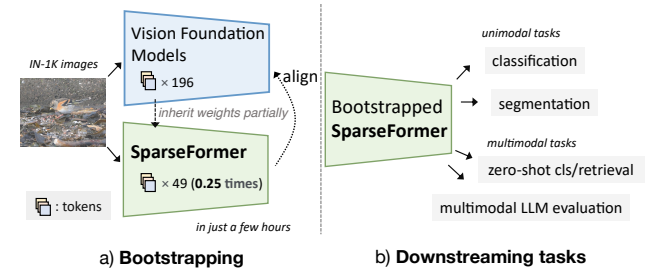


Figure 1. SparseFormer bootstrapping procedure and task evaluation. a) With only images as inputs, we bootstrap SparseFormers from vision foundation models by inheriting weights and aligning final representations with much fewer tokens (e.g., 0.25 \times). b) Bootstrapped SparseFormers can serve as the efficient vision encoder in either off-the-shelf or fine-tuning manner for both unimodal and multimodal tasks.

for its high capacity and proven effectiveness in modeling massive natural language corpora [4, 64]. This large capacity is also observed in the vision domain. For instance, the large-scale ViT pre-trained on the JFT-300M dataset [60] delivers notable performance on various down-streaming tasks and scaling up such models and data still consistently leads to consistent improvements [13, 49, 76]. Despite their improved performance, these vision foundation transformers require significantly more computational resources and memory, both in the training and inference stage, especially when working with high-resolution images, due to the transformer architecture. As an example, processing a single 384×384 resolution image with ViT-L/16 requires handling 576 visual tokens, and attention operators between these tokens take up memory and computation quadratically with the number of tokens.

Recently, SparseFormer [20] has been proposed as an alternative vision transformer architecture with much fewer visual tokens in the latent space rather than the original image space. Each token is associated with a region of interest (RoI) descriptor, and SparseFormer exploits the recurring focusing transformer to adjust token RoIs and sample image features sparsely according to these RoIs. This design allows tokens to be deformable and adjustable in terms of their spatial locations. By training with only classification

✉: Corresponding Author.

labels, these latent tokens can focus on foreground objects and exclude non-informative backgrounds. Therefore, the number of tokens in SparseFormer can be greatly reduced compared to conventional vision transformers, as well as the computational cost and memory footprint. Though being effective in the low compute region, SparseFormer is found to be hard to scale up, and the largest SparseFormer variant in only achieves 84.8% top-1 acc on ImageNet 1K [20]. Moreover, despite tokens being reduced, training SparseFormers from scratch is also time-costing on the ImageNet dataset, *e.g.*, ~ 12 A100 GPU days for a base variant, not to mention the large-scale vision language training. Therefore, it is interesting to explore how to efficiently train SparseFormers to serve as efficient visual transformer alternatives but with capabilities as strong as possible.

In this paper, we present a simple solution to this question: to “bootstrap” SparseFormers from large-scale pre-trained vision foundation models, *e.g.*, AugRegs [59] and CLIPs [49]. By “bootstrapping”, we mean to firstly inherit large-scale pre-trained weights from foundation models into the standard transformer encoder blocks in SparseFormers. And then, we randomly initialize the focusing transformer in SparseFormers and train SparseFormers to align the final representation with foundation models with much fewer tokens, as depicted in Figure 1. For preserving the structure of the output space, we only tune a moderate early blocks with pre-trained weights. Thanks to the SparseFormer efficiency and reuse of pre-trained parameters, we can quickly bootstrap scaled-up SparseFormer variants in few hours. For instance, bootstrapping from AugReg-ViT-L/16 (85.8% on IN-1K with 388 imgs/s) only takes just 6 hours on 8 A5000s, and the resulted SparseFormer achieves 84.5% IN-1K accuracy, using only 49 visual tokens, with 1557 imgs/s throughput. Continuing bootstrapping with 80 tokens leads to 85.5% accuracy, only 0.3% lagging behind AugReg-ViT-L/16 but with $2.4\times$ throughput. Bootstrapped SparseFormers can also serve as backbones for semantic segmentation, reaching 51+ mIoU on ADE20k [78] via 256 tokens for a 512×512 input.

Moreover, since the bootstrapping procedure only needs images as inputs without any labels, we can also bootstrap SparseFormers from CLIP models to output the visual embedding in the language-aligned space just on ImageNet-1K. Without seeing any caption, SparseFormer bootstrapping from CLIP ViT-L/14-336 demonstrates 75.9% ImageNet-1K zero-shot accuracy with only $0.25\times$ tokens, as well as the 57.0% I \rightarrow T@1 retrieval score on the out-of-domain MS COCO [37]. In addition to that, we can also incorporate our bootstrapped SparseFormers into multimodal large language models (MLLMs) seamlessly without further fine-tuning and obtain promising results on the multimodal question answering ScienceQA dataset [41].

2. Related work

2.1. Vision Foundation Models

The term “foundation model” was first introduced in Bommasani et al. [3] to refer to highly capable language transformers that were once pre-trained on massive data samples and can be easily adapted to various downstream tasks. Exemplary of such models includes BERT [15] and GPTs [4, 47, 48]. The concept “foundation model” also applies to computer vision area as well, dating back to the previous cornerstone vision backbones [22, 57] pre-trained on the ImageNet. Recent work [16, 30, 39, 59] on computer vision focuses on pre-training models on large-scale datasets, such as ImageNet-21K [14] and JFT-300M [60], with the classification supervision. The reliance on human annotations of these datasets somewhat imposes a constraint on pre-training such foundation models and several work explores contrastive learning [7, 21, 23], masked modeling [18, 24, 58, 62] or other self-supervised learning [5, 45] to alleviate this issue. Despite the hunger for labels, training time and hardware demands make training a vision foundation model a difficult and expensive thing.

Data in various domains, such as web images [55], is naturally multimodal in language and vision, which invokes interests in pre-training vision models with languages [32, 33, 61, 68, 72, 73, 77]. For example, CLIPs [49] as the most canonical multimodal models align the output embedding spaces of the language transformer and the vision transformer, leading to remarkable zero-shot and transferring capabilities on various vision-language understanding tasks. The success of CLIP has also benefited many downstream tasks, including text-to-image generation [51, 53, 54] and open-vocabulary detection [31, 35, 42, 79]. Furthermore, given the growing influence and availability of large language models (LLMs) [44, 64], recent studies [12, 33, 34, 38, 80] aim to build multimodal models with image inputs by using CLIP pre-trained vision transformers. But still, training these CLIP vision encoders from scratch can be extremely expensive, and such dense transformers can be computationally demanding during inference, particularly for inputs with high resolution.

2.2. Efficient Vision Foundation Models

It is always an appealing topic to build efficient vision transformers, *e.g.*, with efficient attention mechanisms [10, 29, 65], or using compact transformer architectures [17, 25, 39, 66]. Beside these, a research line aims to investigate the token redundancy in vision and thus alleviate it [2, 6, 19, 36, 52, 67, 74] to expedite the inference phase. Knowledge distillation on [1, 26, 43, 63, 69] have also been widely applied to architecting small transformers with knowledge transferred from large vision transformers. Apart from unimodal ones, there has been an increasing em-

phasis on efficient multimodal transformers [46, 56, 70].

The recently proposed SparseFormer [20] as a vision transformer variant exploits the token adjusting mechanism with much fewer tokens to expedite the visual understanding task. SparseFormer can be also seen as a vision transformer with the token reduction right from the start, but also requires lengthy training from scratch. In this paper, we mainly discuss how to bootstrap SparseFormers from vision foundation models with limited training time and data samples, and use the bootstrapped SparseFormers as the efficient vision foundation models with most preserved performance in both unimodal and multimodal settings.

3. Method

In this section, we will first briefly revisit the SparseFormer architecture [20]. Next, we will describe in detail bootstrapping SparseFormers from vision foundation models with the limited training time, hardware budgets, and data samples.

3.1. Prerequisites: the SparseFormer Architecture

The SparseFormer architecture, as a vision transformer variant, aims to represent an image by a highly reduced number of tokens and their corresponding adjustable RoIs. SparseFormer consists up of two components: the focusing transformer and the cortex transformer. The focusing transformer, which is designed to be with minimal parameters and computational costs, iteratively adjusts token RoIs to focus on foregrounds. The cortex transformer is exactly a plain vision transformer encoder, similar to [16], which is with the majority of parameters and computation used by SparseFormer. Specifically, given the latent token embedding set $\mathbf{T} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_N]$ and the token RoI set $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N]$, where N is the number of latent tokens, the focusing transformer operates as follows:

$$\mathbf{T}^{i+1}, \mathbf{B}^{i+1} = \text{FocusingTransformer}(\mathbf{T}^i, \mathbf{B}^i), \quad (1)$$

where \mathbf{T}^1 and \mathbf{B}^1 are initial token embedding set and initial RoI set, and they are learnable parameters of models, independent of input images. In the typical setting, the focusing transformer is repeated in 4 iterations. In each iteration, the focusing transformer extracts image features according the current token RoI $\mathbf{b}_{(\cdot)} = [x, y, w, h]$ into token embedding $\mathbf{t}_{(\cdot)}$, perform self-attention and feed forward network on embeddings, and adjust token RoIs. The RoI adjusting mechanism uses the normalized delta form:

$$\begin{aligned} x' &= x + \Delta_x w, y' = y + \Delta_y h, & (2) \\ w' &= w \cdot \exp(\Delta_w), h' = h \cdot \exp(\Delta_h), & (3) \end{aligned}$$

where the tuple $(\Delta_x, \Delta_y, \Delta_w, \Delta_h)$ is produced by the token embedding in the focusing transformer. All operations are applied individually to each token, except for self-attention which involves token interaction.

After the focusing transformer, these token RoIs tend to cluster around foreground objects [20]. Then, the so-called cortex transformer, whose architecture is exactly a standard transformer encoder, operates on the token embedding as:

$$\mathbf{T}^{\text{final}} = \text{CortexTransformer}(\mathbf{T}'), \quad (4)$$

where we can see that the cortex transformer just processes token embeddings and does not rely on the token RoIs anymore. With the RoI adjusting mechanism, SparseFormer utilizes a highly reduced number of tokens (e.g., $0.25\times$) for visual understanding by only attending to foreground objects and excluding backgrounds.

Feasibility to bootstrap SparseFormers from ViTs.

Since the cortex transformer strictly follows the standard transformer encoder architecture, it is feasible to inherit pre-trained vision transformer weights. Once the weights of the cortex transformer are inherited, we only need to train the focusing transformer from scratch to align its representation into the cortex transformer to be compatible with inherited cortex transformer weights. Thus, when the input representation well aligned, the output representation of SparseFormer will also align with the original pre-trained vision transformers, which is especially beneficial for language-aligned vision representations, such as CLIPs [49]. We term this procedure as “bootstrapping” SparseFormers from large-scale pre-trained vision foundation models.

3.2. Bootstrapping from Vision Foundation Models

As discussed above, we need to align the output representation of the focusing transformer with the input representation of inherited vision transformers. However, this straightforward idea is infeasible since the shape of these two representations typically does not match as the number of SparseFormer tokens is much lower than ones in most vision transformers.

Therefore, we resort to directly aligning the final representation of SparseFormer to one from pre-trained vision transformers. Specifically, we propose a simple method to bootstrap SparseFormers from vision foundation models, which are mostly based on transformers. First, we design SparseFormer variants to match the dimension and the number of their cortex transformer blocks with these pre-trained vision transformers. We inherit these large-scale pre-trained weights of these transformers and load them into cortex transformer blocks in SparseFormers. Then, we initialize the focusing transformer from random and concatenate these two to build a complete SparseFormer.

Once a SparseFormer is constructed, we align the final representation from SparseFormer with that of the pre-trained vision transformer given the same image input, as depicted in Figure 2. We simply use the cosine loss to align

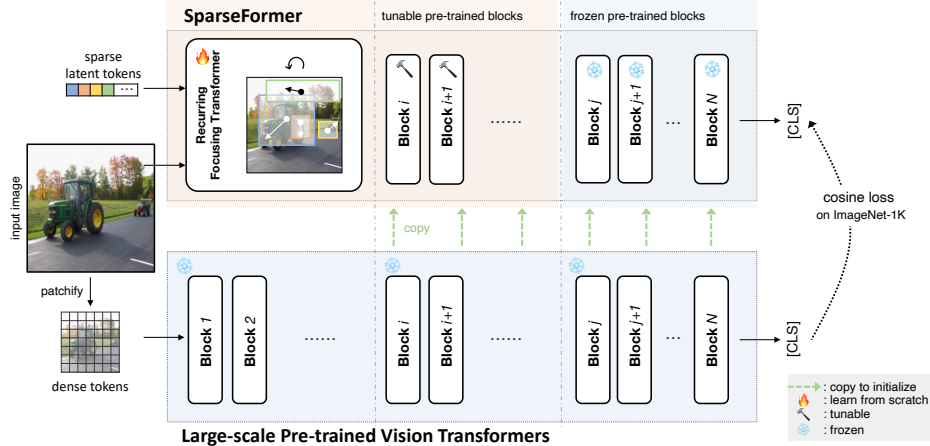


Figure 2. The detailed bootstrapping procedure. We typically set the number of sparse latent tokens in SparseFormers to 1/4 those in vision transformers. The starting index of tunable blocks i is $N/3$ and the frozen is $2N/3$ for all bootstrapping settings. Note that [CLS] represents the extra token in vision transformers besides visual tokens and there is no classification supervision on [CLS] in bootstrapping.

these two final token representations:

$$\ell = 1 - \frac{\mathbf{t}_{\text{sparseformer}}^T \mathbf{t}_{\text{pretrained}}}{\|\mathbf{t}_{\text{sparseformer}}\| \|\mathbf{t}_{\text{pretrained}}\|}, \quad (5)$$

where $\mathbf{t}_{\text{sparseformer}}$ and $\mathbf{t}_{\text{pretrained}}$ are the final token embeddings in the last embedding space of models, which are [CLS] tokens before the classifier layer in classification models or after the projection layer in CLIP models. We can bootstrap SparseFormers from large-scale pre-trained vision foundation models without labels, since the alignment target from pre-trained models, $\mathbf{t}_{\text{pretrained}}$, already has rich semantics. All we need to do is to seek for a set of unlabeled supporting images. We find that the medium-sized ImageNet-1K training set [14], which is most publicly available to its scale, with minimal augmentation is a sufficient supporting image set for bootstrapping competitive SparseFormers from pre-trained models.

Difference with distillation methods. It is important to note that the cosine loss between two token embeddings is not dependent on pre-defined label sets, and we do not exploit any labels or captions during the bootstrapping process. This differs a lot from existing distillation methods, which typically make use of specified label sets [26, 63] for classification models or vision-language pairs for multi-modal models [70]. Also, distillation methods do not reuse teacher model weights since the student architecture usually differs from the teacher, and thus require a lengthy schedule to transfer knowledge effectively. In comparison, our method can bootstrap a 200M SparseFormer variant in a few hours (20 epochs) with only 1.2M images from ImageNet as inputs.

Truncate the leading, tune the middle, and freeze the ending blocks. The leading blocks of pre-trained vision transformers are usually specialized for low-level visual modeling [50], which overlaps with the role of the focusing transformer in SparseFormers. Thus, we choose to discard a moderate number, 1/3, of leading transformer blocks when constructing our SparseFormer variants, which also help reduce the compute need. Since the bootstrapping goal is to well align the final representation, we set a few middle transformer blocks tunable starting from inherited weights to adapt the output of the focusing transformer. We leave frozen the rest transformer blocks in the ending to preserve the structure of the output space by pre-trained weights as much as possible.

4. Experiments

We start experiments on bootstrapping SparseFormers from large-scale pre-trained unimodal models and then discuss vision-language pre-trained ones.

4.1. Bootstrapping from Unimodal Models

We choose the well-established AugReg vision transformers [59], which strictly follow the original ViT architecture, as the unimodal classification models to bootstrap from. The AugReg models are pre-trained by the supervised classification on the ImageNet-21K dataset with well-curated data augmentation and regularization techniques.

Model configurations and experimental settings. We design two SparseFormer variants (SF-B_{AugReg} and SF-L_{AugReg}) according to the base and large variant (ViT-B/16-224_{AugReg} and ViT-L/16-224_{AugReg}) in Table 1. We set the token dimension in the focusing transformers half those of

the cortex transformers except that the last focusing transformer block performs sampling operations on the full dimension, in accord with [20]. Unlike the leading one, the last focusing transformer block does not involve FFNs and self-attention layers. We append a learnable [CLS] to latent tokens after the focusing stage in SparseFormer and use its final representation for bootstrapping and classification. To be more consistent with AugReg models (CLIPs below as well) which add positional encoding to visual tokens, we inject positional encoding based on token RoIs in the sinusoidal form in each focusing iteration. To further reduce compute and make SparseFormer even more sparse, we reduce the channel of the early convolution from 96 to 64, and shrink the number of sampling points for a token in the sparse sampling procedure from 36 to 16.

model	#tokens	focusing iterations	focusing dimension	#truncate blocks	#unable blocks	FLOPs	#Params
ViT-B/16-224 _{AugReg}	196	-	-	-	-	17.5G	86M
SF-B _{AugReg}	49	4+1	384	4	4	3.8G	64M
ViT-L/16-224 _{AugReg}	196	-	-	-	-	61.6G	304M
SF-L _{AugReg}	49	4+1	512	8	8	11.4G	213M

Table 1. Configurations of SparseFormers.

We choose the publicly-available medium-sized ImageNet-1K training set as our bootstrapping image set. The data augmentation is set to be minimal: random horizontal flipping with the probability 0.5 and random resized cropping with from the scale (0.5, 1.0) and the aspect ratio (3/4, 4/3). We use a budget of 20 epochs as our bootstrapping schedule, where learning rate begins at 2×10^{-4} and follows a half-cosine decay schedule after the first warming-up epoch. In the warming-up epoch, the inherited pre-trained weights keep frozen for training stability. In bootstrapping, only the focusing transformer and tunable blocks of cortex transformer are learnable. All the rest parameters inherited from AugReg models, including the final layer normalization and classification layer, remain frozen.

Bootstrapped SparseFormers with 49 tokens. We show results of bootstrapped SparseFormer from AugReg classification models in Table 2. The default token number is 49 for SparseFormer models. From the results, we can see that SparseFormer models with 49 tokens have achieved promising results, *e.g.*, 84.9% top-1 accuracy for SF-L_{AugReg} bootstrapping from ViT-L/16-384_{AugReg}. Note that the number of tokens, *e.g.* 49, is fixed from the beginning to the end in SparseFormer models, which differs from recent other methods progressively dropping tokens through blocks [2, 6]. To our best knowledge, we are the first to obtain $\sim 85\%$ top-1 accuracy on ImageNet-1K with only 49 tokens in all blocks for vision transformers.

model	top-1 acc.	#Params	FLOPs	img/s
ViT-B/16-224 _{AugReg} , 196 tokens	84.6	86M	17.5G	1126
SF-B _{AugReg} , 49 tokens	82.5	64M	3.8G	3001 $\uparrow 166\%$
SF-B _{AugReg} , 80 tokens \uparrow	83.4	64M	6.2G	2080 $\uparrow 85\%$
SF-B _{AugReg} , 80 tokens \uparrow , 320px \uparrow	83.7	64M	6.3G	1898 $\uparrow 67\%$
ViT-L/16-224 _{AugReg} , 196 tokens	85.8	304M	61.6G	388
SF-L _{AugReg} , 49 tokens	84.5	213M	11.4G	1557 $\uparrow 301\%$
SF-L _{AugReg} , 80 tokens \uparrow	85.2	213M	18.6G	1008 $\uparrow 160\%$
SF-L _{AugReg} , 80 tokens \uparrow , 320px \uparrow	85.5	213M	18.7G	928 $\uparrow 139\%$
ViT-L/16-384 _{AugReg} , 576 tokens	86.9	304M	191G	105
SF-L _{AugReg} , 49 tokens	84.9	213M	11.6G	1353 $\uparrow 1189\%$
SF-L _{AugReg} , 144 tokens \uparrow	86.5	213M	33.9G	560 $\uparrow 433\%$
SF-L _{AugReg} , 196 tokens \uparrow	86.7	213M	46.5G	378 $\uparrow 260\%$

Table 2. Results of SparseFormers bootstrapped from AugReg models. The *default* SparseFormer token number is 49. Throughputs are measured with half-precision data type and batch size of 128 on one A5000 GPU.

Continue to bootstrapping with more tokens. Although SparseFormers with 49 tokens yield promising results, they are still lagging behind those models to bootstrap from with a non-negligible margin. To minimize the gap, we further continue to bootstrapping SF models with the increased number of tokens from default SF models. Note that these increased numbers are still much lower than original vision transformers, especially for high resolution one. The continued bootstrap lasts for 5 epochs with LR starting from 5×10^{-5} , also following a half-cosine decay scheme. We denote these models with an extra notation \uparrow . From the table, the gaps between SparseFormers and AugRegs become closer when original models become larger, especially for SF-L_{AugReg} with 196 tokens, which is on par with ViT-L/16-384_{AugReg} but at about $3 \times$ real-time throughput. When combined with the higher input resolution, SparseFormer models still yield improvements on the same number of tokens at a marginal additional burden. Note that as a sparse architecture, the input resolution has little impact on computational costs of the most components in SparseFormers.

Ablation studies on bootstrapping. Our bootstrapping procedure is simple as it just aligns final representations of SparseFormers and large-scale pre-trained models given the same image input. Labels are not needed through the bootstrapping procedure except for the evaluation on bench-

bootstrapping	with	w/o reusing	KL distill	KL distill	supervised cls from scratch
	cls loss	weights	$\tau = 3$	$\tau = 5$	
82.5	77.4	69.5	81.0	81.5	80.0

Table 3. Ablations on bootstrapping SF-B_{AugReg} with 49 tokens. Except for ‘from scratch’ (supervised classification, from scratch, 300ep), we keep the bootstrapping training settings the same (*e.g.*, trainable parameters, epochs & augmentation).

model	top-1 acc.	FLOPs	img/s
ViT-B/16-224 _{AugReg}	84.6	17.5G	1126
ToMe@16, off-the-shelf	80.4	8.8G _{↓49.7%}	1980 _{↑75%}
ToMe@16, fine-tuned	82.8	8.8G _{↓49.7%}	1980 _{↑75%}
DiffRate, fine-tuned [†]	83.2	11.5G _{↓34.3%}	1400 _{↑24%}
SF-B _{AugReg} , 80 tokens [†]	83.4	6.2G _{↓64.6%}	2080 _{↑84%}
SF-B _{AugReg} , 80 tokens [†] , 320px [†]	83.7	6.3G _{↓64.0%}	1898 _{↑68%}
SF-B _{AugReg} , 100 tokens [†] , 320px [†]	84.0	7.8G _{↓55.4%}	1579 _{↑40%}
ViT-L/16-224 _{AugReg}	85.8	61.6G	388
ToMe@12, off-the-shelf	20.6 [‡]	20.8G _{↓66.2%}	1006 _{↑159%}
ToMe@8, off-the-shelf	83.5	31.0G _{↓49.7%}	700 _{↑80%}
ToMe@8, fine-tuned	84.8	31.0G _{↓49.7%}	700 _{↑80%}
DiffRate, fine-tuned [†]	85.7	42.3G _{↓31.3%}	452 _{↑16%}
SF-L _{AugReg} , 80 tokens [†]	85.2	18.6G _{↓69.8%}	1008 _{↑160%}
SF-L _{AugReg} , 80 tokens [†] , 320px [†]	85.5	18.7G _{↓69.6%}	928 _{↑139%}
SF-L _{AugReg} , 100 tokens [†] , 320px [†]	85.7	23.4G _{↓62.0%}	752 _{↑94%}
ViT-L/16-384 _{AugReg}	86.9	191G	105
ToMe@40, off-the-shelf	8.8 [‡]	56.0G _{↓70.7%}	315 _{↑200%}
ToMe@23, off-the-shelf	86.2	96.4G _{↓49.5%}	187 _{↑78%}
ToMe@23, fine-tuned	86.8	96.4G _{↓49.5%}	187 _{↑78%}
SF-L _{AugReg} , 196 tokens [†]	86.7	46.5G _{↓75.7%}	378 _{↑260%}

Table 4. Comparisons on SparseFormers with token reduction methods, ToMe [2] and DiffRate [6]. [†] use MAE [24] pre-trained weights. [‡] ToMe is not unfriendly to large token merging rates.

marks. Here we ablate the bootstrapping design in Table 3. Interestingly, aligning the final representations along with the classification loss (‘with cls loss’ in the table) actually impairs the performance. We suspect that the supervised classification requires strong augmentation and regularization to be stable with the alignment. We also replace our token representation alignment with KL distillation loss with different temperatures on the ImageNet labels. This simple distillation yields promising results but still with $> 1\%$ lagging behind our bootstrapping.

We also include the bootstrapping result without the reuse of pre-trained weights (but the final linear classifier remained) and train the entire transformer to align the final representation for 20 epochs. As expected, it yields the inferior result, and more training epochs may be required to well align from scratch.

Comparison with token reduction methods. SparseFormers can be considered as a method to reduce the number of visual tokens in vision transformers after pre-training, similar to token pruning or merging approaches [2, 6, 19, 36]. Here, we investigate the effectiveness of SparseFormers from the token reduction perspective and compare them to state-of-the-art reduction methods, ToMe [2] and DiffRate [6], in Table 4. In addition to using off-the-shelf AugReg models with ToMe, we also fine-tune AugRegs with ToMe using the fine-tuning recipe described in [24] to ensure a fair comparison with our methods, as additional fine-tuning is required. We sweep to find best fine-

model	mIoU	#Params	FLOPs	train mem.	img/s
Swin-L [39] + UPerNet [71] [†]	51.1	234M	408G	11G, 2bs	20
ViT-Adapter-B _{AugReg} [8] + UPerNet [71]	51.9	134M	376G	50G, 2bs	19
ViT-Adapter-L _{AugReg} [8] + UPerNet [71]	53.4	364M	667G	49G, 2bs [‡]	12
SF-B _{AugReg} segmentation (ours)	49.3	68M	36G	10G, 4bs	88
SF-L _{AugReg} segmentation (ours)	51.5	216M	77G	13G, 4bs	46

Table 5. Segmentation performance on ADE20k val set [78]. [†] is reported by mmsegmentation [11]. [‡] is with gradient checkpointing. Throughput is measured with batch size 4 on A5000.

tuning LRs for these ToMe AugRegs respectively. DiffRate on MAE models are also included, where the ViT-L/16-224_{MAE} baseline reports a slightly better result than ViT-L/16-224_{AugReg}. As shown in the table, SparseFormer models achieve the best trade-off between the accuracy and the actual throughput, especially for the large models where many tokens may be redundant. In fact, tokens in bootstrapped SparseFormers can be further applied with these token reduction methods but this is beyond this paper scope.

Dense prediction task. We here further investigate the ability of bootstrapped SparseFormers as pre-trained backbones to perform dense per-pixel prediction task, *e.g.*, semantic segmentation. As latent tokens of SparseFormer are not structured in a grid-like map like conventional vision transformers [16, 39, 66], it requires some workarounds to obtain dense feature maps from these latent tokens. We generally follow the idea of the original SparseFormer on semantic segmentation to project latent tokens back into the original dense pixel space according to their RoIs [20]. Here, we find that it is more effective to straightforward project prediction logits of tokens, instead of embeddings, back into the dense pixel space using simple single-head attention aggregation:

$$\mathbf{P}_{\text{token}} = \text{classifier}(\mathbf{T}) \in \mathbb{R}^{N \times L}, \quad (6)$$

$$\mathbf{P}_{\text{dense}} = \text{softmax}(\mathbf{Q}_{\text{dense}} \mathbf{K}_{\text{token}}^T / \sqrt{d} + \mathbf{B}) \mathbf{P}_{\text{token}}, \quad (7)$$

where $\text{classifier}(\cdot)$ just classify the latent token into prediction logits, N is the number of latent tokens, and L is the number of classes. In the current implementation, we first use a linear layer to reduce the latent token dimension to $d = 256$ and introduce two small transformer encoder blocks ($d = 256$) to process latent tokens, and we use a linear layer as our latent classifier. $\mathbf{Q}_{\text{dense}} \in \mathbb{R}^{HW/4^2 \times d}$ is the flattened feature map transformed by one 3×3 convolution from the early convolution feature in SparseFormer with shape $H/4 \times W/4$. We bias the attention score by $\mathbf{B} = \mathbf{B}_{\text{geometric}} + \mathbf{B}_{\text{predictive}}$, where the geometric bias $\mathbf{B}_{\text{geometric}}$ is in the Gaussian-like form w.r.t. token RoIs as described in [20]. The predictive bias is produced by a linear layer whose input is the latent token embedding, $\mathbf{B}_{\text{predictive}} = \text{linear}(\mathbf{T}) \in \mathbb{R}^{1 \times N}$, indicating the overall

method	visual encoder	visual FLOPs	IN-1K zero-shot@1	MS COCO I→T@1	MS COCO T→I@1	Flickr30k I→T@1	Flickr30k T→I@1
CLIP	ViT-B/16-224	17.5G	68.3	52.4	33.0	81.9	62.0
CLIP	ViT-B/32-224	4.4G	63.3	50.1	30.4	77.5	58.8
TinyCLIP [70]†	ViT-63M/32-224	2.0G	64.5	56.9	38.5	84.9	66.0
CLIP + SF (ours)	SF-B _{CLIP} w/ 49 tokens	3.8G	66.0	47.7	31.5	76.3	59.6
CLIP	ViT-L/14-224	81.1G	75.4	56.3	36.5	85.1	65.2
CLIP + SF (ours)	SF-L _{CLIP} w/ 64 tokens	14.8G	73.6	52.3	34.8	80.3	62.8
CLIP + SF (ours)	SF-L _{CLIP} w/ 128 tokens	30.0G	74.9	54.7	36.4	82.4	65.2
CLIP	ViT-L/14-336	191.0G	76.6	57.9	37.1	87.4	67.3
CLIP + SF (ours)	SF-L _{CLIP} w/ 144 tokens	33.9G	75.9	57.0	38.3	84.6	67.7

Table 6. Zero-shot and retrieval evaluation results. We individually bootstrap SparseFormers from CLIP visual encoders in the grayed rows with <2M ImageNet-1K images as inputs. † needs training on text-image pairs of LAION-400M [55].

importance of token logits into the dense map. In simple words, we map the latent token predictions into the dense grid with the consideration of their RoIs and semantic importance, similar to [9]. The classification supervision is imposed on the mapped dense predictions.

We generally follow the training recipe of [8], including learning rate and drop path [27] settings, and finetune our bootstrapped SparseFormers (default ones w/ 49 tokens) on the ADE20k dataset [78] with the increased token number of 256. As our segmentation models based on SparseFormer are memory friendly, the batch size is increased to 4 and the total training iteration is reduced to 80K. The results are reported in Table 5. We can see that the performance of SparseFormer on the segmentation task is promising, with much fewer FLOPs and much higher throughput. This efficiency is credited to fewer computation over limited latent tokens in SparseFormers compared to conventional dense vision transformers (*e.g.*, 256 versus 1024 for ViT-Adapter [8] for 512×512 inputs).

4.2. Bootstrapping from Multimodal Models

Till now, we have bootstrapped SparseFormers from unimodal classification models and apply them to the downstream dense prediction task, showing their efficiency and effectiveness. Since our bootstrap procedure directly aligns the final representations to the pre-trained transformers, it can be easily adopted for the fundamental vision-language CLIP [49] models as well.

Experimental setups. The bootstrapping procedure for CLIP models exactly follows the same recipe in unimodal classification AugRegs, as described in Sec 4.1, including the training budget, the data augmentation, and the learning rate setting. We use the official OpenAI CLIP pre-trained vision transformers to bootstrap from, *i.e.*, ViT-B/16-224, ViT-L/14-224, and ViT-L/14-336. We only bootstrap SparseFormers from the CLIP visual encoders and leave the text encoders untouched. It is worth noting that we still only use the ImageNet-1K training set as the sup-

porting image dataset, and we do not leverage any text-image pairs to align the SparseFormer image embeddings with CLIP models. This differs from TinyCLIP [70], which also inherits pre-trained weights but requires at least 15M text-image pairs to distill smaller models.

Results. We use the frozen CLIP text encoders in combination with our bootstrapped SparseFormers as visual encoders to perform the zero-shot classification task on ImageNet-1K validation set, as well as the retrieval task on MS COCO val set [37] and the Flickr30k benchmark [75]. Results are presented in Table 6. With only images as samples for aligning in the bootstrapped procedure, SparseFormers as the visual encoders achieve decent zero-shot accuracies with greatly reduced compute, especially for large and high resolution models. For instance, our SparseFormer-L with 144 tokens is just 0.7 point behind the ViT-L/14-336 with 576 tokens, and surpass the ViT-L/14-224 by 0.5 margin with less than 0.5× computational cost. Interestingly, the zero-shot retrieval performance on MS COCO and Flickr30k, which have more complicated data distributions compared to the ImageNet samples for aligning, is also quite effective.

4.3. Multimodal Large Language Models

Building large language models with integrated vision capabilities is an emerging and appealing topic [33, 34, 38] within the computer vision community. Since we have bootstrapped SparseFormers from CLIP models, a natural question arises: can we seamlessly incorporate SparseFormer as the vision encoder in the multimodal large language model

vision encoder	acc	multi-modal only acc
ViT-L/16-224	91.2	88.6
SF-L _{CLIP} w/ 64 tokens	89.2	84.5
SF-L _{CLIP} w/ 128 tokens	89.8	85.8

Table 7. In-place vision encoder replacement in LLaVa on the ScienceQA test set.

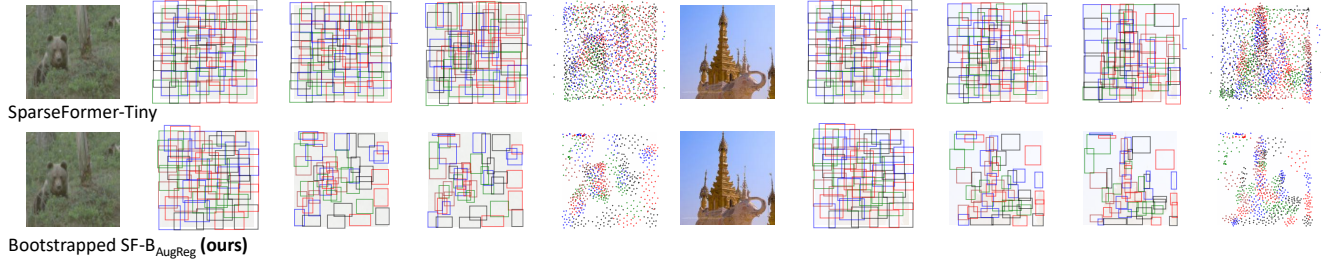


Figure 3. Visualization on the original SparseFormer-Tiny [20] and our bootstrapped ViT-B/16-224_{AugReg}. For each image, there are an input image, token RoIs in the {first, third, last} stage, and sampling points in the last stage in the focusing transformer from left to right.

without re-finetuning LLMs. The answer is yes. We experiment with LLaVa [38], a recent large language model that is fine-tuned for instructions and can accept images and texts as the prompt. The vision encoder in LLaVA is typically the CLIP ViT-L and it directly outputs all visual tokens (typical number is 256 for a 224² image) as input tokens for autoregressive language modeling, meaning that there are already 256 “word tokens” for image modeling before user input prompt. In comparison, SparseFormers reduce this number of tokens to typically 1/4 \times , which is more friendly for LLM text modeling.

To demonstrate quantitative multimodal LLM results, we benchmark the fine-tuned LLaVa variant, `llava-1cs558k-scienceqa-vicuna-13b-v1.3`, for the instruction following on the ScienceQA multi-choice question answering task [41]. We find that the data distribution is unique on the ScienceQA dataset and directly using bootstrapped SparseFormers in Sec 4.2 leads to inferior results. Therefore, we continue bootstrapping SparseFormers from CLIP ViT-L/14-224 using the ScienceQA training images for ~ 1000 iterations to adapt to this image domain for about 20 minutes. Then we replace the CLIP vision encoder in LLaVa with SparseFormers without any further fine-tuning LLM. Results are shown in Table 7. The bootstrapped SparseFormer as the vision encoder preserves most language-vision abilities (84.5 vs 88.6 in the multimodal only entry) with only 1/4 \times visual tokens as inputs to the following autoregressive LLM. The performance gap between SparseFormers and CLIP even narrows further when scaling to 128 tokens, which is still half the size of CLIP tokens.

4.4. Visualization

We show the visualization of our bootstrapped SparseFormer and compare it to the original SparseFormer [20] in Figure 3. While the bootstrapped ViT-B/16-224_{AugReg} and the original SparseFormer-Tiny use exactly the same number of visual tokens (*i.e.*, 49), our bootstrapped model demonstrates the better capability to focus on foregrounds and exclude non-informative regions with less sampling points (36 \rightarrow 16 per token, as described in Section 4.1). As a

result, the bootstrapped SparseFormer from large-scale pre-trained models can be *even sparser* in both focal ability and sampling efficiency compared with the original one.

5. Conclusion

In this paper, we have proposed a simple but effective method to bootstrap SparseFormers from large-scale pre-trained vision foundation transformers by inheriting most pre-trained weights and explicitly aligning the final representations. With the short training time and limited data samples, the SparseFormers bootstrapped from unimodal classification models can maintain the pre-trained capabilities as much as possible with much fewer tokens and higher real-time throughput. In particular, we can bootstrap a SparseFormer with only 49 tokens which obtains 84.9 top-1 accuracy on ImageNet 1K. Also, SparseFormers can serve as backbone networks in the dense per-pixel semantic segmentation task with decent results but at 2 \times throughput. In addition to classification models, SparseFormers can be bootstrapped from multimodal pre-trained transformer CLIPs without seeing any text caption in the bootstrapping procedure and the bootstrapped SparseFormers demonstrate strong multimodal performance with down to 0.2 \times FLOPs.

Limitation The limitation of our proposed bootstrapping procedure is that it presumes the underlying transformer architecture of vision foundation models and thus can only applied to transformer-based foundation models. This is mostly the case, but there are some exceptions, *e.g.*, [40]. Also, it necessitates the availability of weights of vision foundation models, while that may not be possible for some large-scaled proprietary ones.

Acknowledgement

This project is supported by the National Research Foundation, Singapore under its NRFF Award NRF-NRFF13-2021-0008.

References

- [1] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10925–10934, 2022. 2
- [2] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your ViT but faster. In *International Conference on Learning Representations*, 2023. 2, 5, 6
- [3] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 2
- [4] 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. *NeurIPS*, pp. 1877–1901, 2020. 1, 2
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pp. 9630–9640. IEEE, 2021. 2
- [6] Mengzhao Chen, Wenqi Shao, Peng Xu, Mingbao Lin, Kaipeng Zhang, Fei Chao, Rongrong Ji, Yu Qiao, and Ping Luo. Diffrate : Differentiable compression rate for efficient vision transformers. *CoRR*, abs/2305.17997, 2023. 2, 5, 6
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020. 2
- [8] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. In *ICLR*. OpenReview.net, 2023. 6, 7
- [9] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. 2021. 7
- [10] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarnlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020. 2
- [11] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 6
- [12] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023. 2
- [13] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. In *International Conference on Machine Learning*, pp. 7480–7512, 2023. 1
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009. 2, 4
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1, 2
- [16] 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. 1, 2, 3, 6
- [17] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. *arXiv:2104.11227*, 2021. 2
- [18] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. EVA: exploring the limits of masked visual representation learning at scale. In *CVPR*, pp. 19358–19369. IEEE, 2023. 2
- [19] Mohsen Fayyaz, Soroush Abbasi Koohpayegani, Farnoush Rezaei Jafari, Sunando Sengupta, Hamid Reza Vaezi Joze, Eric Sommerlade, Hamed Pirsiavash, and Jürgen Gall. Adaptive token sampling for efficient vision transformers. In *European Conference on Computer Vision*, pp. 396–414, 2022. 2, 6
- [20] Ziteng Gao, Zhan Tong, Limin Wang, and Mike Zheng Shou. Sparseformer: Sparse visual recognition via limited latent tokens. *arXiv preprint arXiv:2304.03768*, 2023. 1, 2, 3, 5, 6, 8
- [21] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach

- to self-supervised learning. In *Advances in Neural Information Processing Systems*, 2020. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. doi: 10.1109/CVPR.2016.90. 2
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [24] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 6
- [25] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. *arXiv preprint arXiv:2103.16302*, 2021. 2
- [26] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2, 4
- [27] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *ECCV (4)*, volume 9908 of *Lecture Notes in Computer Science*, pp. 646–661. Springer, 2016. 7
- [28] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pp. 4904–4916, 2021. 1
- [29] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020. 2
- [30] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *European Conference on Computer Vision*, 2020. 2
- [31] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. Open-vocabulary object detection upon frozen vision and language models. In *International Conference on Learning Representations*, 2023. 2
- [32] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *Advances in Neural Information Processing Systems*, 2021. 2
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, pp. 12888–12900, 2022. 2, 7
- [34] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2, 7
- [35] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [36] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *International Conference on Learning Representations*, 2022. 2, 6
- [37] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV (5)*, volume 8693 of *Lecture Notes in Computer Science*, pp. 740–755. Springer, 2014. 2, 7
- [38] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 2, 7, 8
- [39] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 6
- [40] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 8
- [41] Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 2, 8
- [42] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European Conference on Computer Vision*, 2022. 2
- [43] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via

- teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 5191–5198, 2020. 2
- [44] OpenAI. Gpt-4 technical report, 2023. 2
- [45] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael G. Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *CoRR*, abs/2304.07193, 2023. 2
- [46] Renjing Pei, Jianzhuang Liu, Weimian Li, Bin Shao, Songcen Xu, Peng Dai, Juwei Lu, and Youliang Yan. Clipping: Distilling clip-based models with a student base for video-language retrieval. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [47] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2
- [48] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, pp. 9, 2019. 2
- [49] Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2, 3, 7
- [50] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, pp. 12116–12128, 2021. 4
- [51] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [52] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *Advances in Neural Information Processing Systems*, 2021. 2
- [53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [54] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 2022. 2
- [55] 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. *CoRR*, abs/2111.02114, 2021. 2, 7
- [56] Dachuan Shi, Chaofan Tao, Ying Jin, Zhendong Yang, Chun Yuan, and Jiaqi Wang. UPop: Unified and progressive pruning for compressing vision-language transformers. In *International Conference on Machine Learning*, 2023. 3
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2
- [58] Mannat Singh, Quentin Duval, Kalyan Vasudev Alwala, Haoqi Fan, Vaibhav Aggarwal, Aaron Adcock, Armand Joulin, Piotr Dollár, Christoph Feichtenhofer, Ross Girshick, et al. The effectiveness of mae pre-pretraining for billion-scale pretraining. *arXiv preprint arXiv:2303.13496*, 2023. 2
- [59] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *Trans. Mach. Learn. Res.*, 2022, 2022. 1, 2, 4
- [60] 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 (ICCV)*, Oct 2017. 1, 2
- [61] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. EVA-CLIP: improved training techniques for CLIP at scale. *CoRR*, abs/2303.15389, 2023. 2
- [62] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in Neural Information Processing Systems*, 2022. 2
- [63] 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*, pp. 10347–10357. PMLR, 2021. 2, 4

- [64] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 2
- [65] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. 2
- [66] 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 *IEEE/CVF International Conference on Computer Vision*, 2021. 2, 6
- [67] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic vision transformers with adaptive sequence length. *arXiv preprint arXiv:2105.15075*, 2021. 2
- [68] Yixuan Wei, Yue Cao, Zheng Zhang, Zhuliang Yao, Zhenda Xie, Han Hu, and Baining Guo. icar: Bridging image classification and image-text alignment for visual recognition. In *Advances in Neural Information Processing Systems*, 2022. 2
- [69] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Tinyvit: Fast pretraining distillation for small vision transformers. In *European Conference on Computer Vision*, pp. 68–85. Springer, 2022. 2
- [70] Kan Wu, Houwen Peng, Zhenghong Zhou, Bin Xiao, Mengchen Liu, Lu Yuan, Hong Xuan, Michael Valenzuela, Xi (Stephen) Chen, Xinggang Wang, Hongyang Chao, and Han Hu. Tinyclip: Clip distillation via affinity mimicking and weight inheritance. In *IEEE/CVF International Conference on Computer Vision*, 2023. 3, 4, 7
- [71] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV (5)*, volume 11209 of *Lecture Notes in Computer Science*, pp. 432–448. Springer, 2018. 6
- [72] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [73] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. 2022. 2
- [74] Hongxu Yin, Arash Vahdat, Jose M. Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. A-vit: Adaptive tokens for efficient vision transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10799–10808, 2022. 2
- [75] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguistics*, 2:67–78, 2014. 7
- [76] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12104–12113, 2022. 1
- [77] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Region-clip: Region-based language-image pretraining. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [78] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *CVPR*, pp. 5122–5130. IEEE Computer Society, 2017. 2, 6, 7
- [79] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, 2022. 2
- [80] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Mini-gpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 2