

Your Image is My Video: Reshaping the Receptive Field via Image-To-Video Differentiable AutoAugmentation and Fusion

Sofia Casarin¹, Cynthia I. Ugwu¹, Sergio Escalera^{2,3}, Oswald Lanz¹

¹Free University of Bozen-Bolzano, Bolzano, Italy

²Computer Vision Center, Barcelona, Spain

³Universitat de Barcelona, Barcelona, Spain

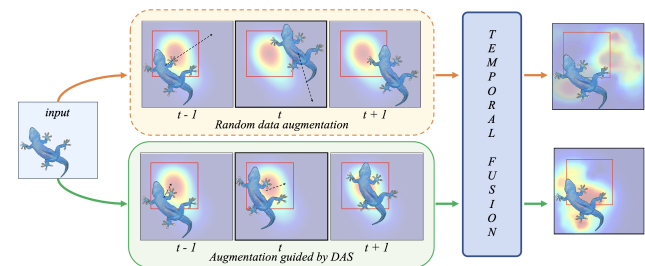
{scasarin, cugwu, olanz}@unibz.it, sergio@maia.ub.es

Abstract


The landscape of deep learning research is moving towards innovative strategies to harness the true potential of data. Traditionally, emphasis has been on scaling model architectures, resulting in large and complex neural networks, which can be difficult to train with limited computational resources. However, independently of the model size, data quality (i.e. amount and variability) is still a major factor that affects model generalization. In this work, we propose a novel technique to exploit available data through the use of automatic data augmentation for the tasks of image classification and semantic segmentation. We introduce the first Differentiable Augmentation Search method (DAS) to generate variations of images that can be processed as videos. Compared to previous approaches, DAS is extremely fast and flexible, allowing the search on very large search spaces in less than a GPU day. Our intuition is that the increased receptive field in the temporal dimension provided by DAS could lead to benefits also to the spatial receptive field. More specifically, we leverage DAS to guide the reshaping of the spatial receptive field by selecting task-dependant transformations. As a result, compared to standard augmentation alternatives, we improve in terms of accuracy on ImageNet, Cifar10, Cifar100, Tiny-ImageNet, Pascal-VOC-2012 and CityScapes datasets when plugging-in our DAS over different light-weight video backbones.

1. Introduction

Creating models with significantly increased capacity, in an attempt to achieve incremental performance improvements, has been the prevailing approach in designing Convolutional Neural Network (CNN) classifiers. As a result, CNNs with increased depth [16, 40, 45, 60], and Vision Transformers (ViTs) [13] were proposed over the years. Specifically, ViT has demonstrated promising results on a wide variety of computer vision tasks including image classifi-



(a) Conceptual representation of the proposed approach. We reshape the Receptive Field (RF) by applying affine transformations optimized through our Differentiable Augmentation Search (DAS). On the top right you can see how fusing with random transformation would not lead to benefits as, when concatenating in time, the employed shift mechanism would fuse features related to random parts. On the bottom, the augmentations guided by DAS obtain specific shapes of the RF so that more context is kept.



(b) RF visualization (ResNet-50, with GSF fusion) when different single or composed transformations are applied. The last column shows our DAS selected operation for CIFAR-10 and CIFAR-100, which combines translation, rotation and zoom. More details in Tab. 4.

Figure 1. **1a** overviews our approach and **1b** shows a real example of obtained receptive fields. The employed transformations are fundamental to shape the receptive field, as shown in **1b**. The augmented images with DAS (Sec. 3.1) are concatenated in time, and processed through a video network that partially shifts and fuses the features (Sec. 3.2).

cation and semantic segmentation [13, 36, 48, 56]. However, while these techniques have demonstrated remarkable success, it is noteworthy that these high-capacity models necessitate increased computational resources for effective training and inference, making them economically impractical for training and deployment within practical application scenarios. Moreover, the over-parametrization of ViT and Deep CNNs makes the networks prone to overfitting,

thus requiring strong regularization to achieve ideal performance. As a result, one alternative trend relies in exploiting the power of data. In this work, we tackle this problem and propose a novel way of augmenting data with the goal of expanding the receptive field of CNNs.

Data augmentation techniques, usually employed to enhance the generalization capabilities of machine learning models, rely on meticulous design, necessitating domain-specific knowledge. To this aim, various auto data augmentation methods, inspired by Neural Architecture Search (NAS), have been extensively utilized in the field of supervised learning [11, 12, 27, 34] to search for accurate augmentation strategies. However, many of these approaches, either require much time [11, 17, 47], forcing the search to be performed on a proxy task or do not propose an effective search strategy, but rather a well defined search space as in [12]. The former, makes a strong assumption that the proxy task provides a predictive indication of the larger task [12]. The dramatic reduction in parameter space of the latter, which allows simple grid search, implies a strong knowledge in the search space definition. As a result, as we empirically show, including “noisy” transformations would imply much larger searching time to achieve the same performance. By harnessing affine transformations optimized through our new Differentiable Augmentation Search (DAS) strategy, we generate a series of images that exhibit motion within a specific region, treating them analogously to video sequences. Inspired by transformation-based models [49] that, operating in the space of affine transforms, rephrase the frames prediction problem as modelling transformation within frames, we introduce a novel methodology to extend the Receptive Field (Fig. 1a). With the underlying hypothesis that augmenting the temporal dimension’s receptive field (RF) could potentially yield advantages for the spatial RF as well, we establish the viability of this approach and empirically showcase its merits by employing Video Networks for processing such image sequences. In this context, to alleviate the unnecessary increase of computational overhead for image-classification and semantic segmentation tasks, we exploit a feature shift mechanism [42], which in the context of video action recognition tasks demonstrates comparable performance to a 3D CNN while keeping 2D CNN complexity. In summary, our contributions are:

- We re-formulate the automatic data augmentation field in a differentiable manner, and propose DAS. By defining a continuous search space of image transformations and exploiting a perturbation-based approach for the transformation selection, we provide a very general and easy-to-deploy alternative to slower existing reinforcement learning methods, and to more search space definition sensitive random ones.
- We propose a new way of handling 2D data by repeat-

edly transforming and concatenating images as frames in a video, obtaining a new perspective to exploit the richness of data. To this aim, we address the question of how the increase of the receptive field in a third dimension impacts the original 2D spatial receptive field.

- We successfully expand, as a result, the receptive field for image classification and segmentation tasks. This allows obtaining ResNet-152 state-of-the-art results for ImageNet while employing a ResNet-50 temporal expanded network, having less than half-parameters, and surpassing DeepLabv3 by 1.3 % on Pascal-VOC and ResNest models by 1.1% on CityScapes datasets.

2. Related Work

In our work, we are interested in image classification and semantic image segmentation tasks. A large focus of the computer vision community has been on engineering better network architectures to improve performance. In earlier designs ImageNet progress was dominated by CNNs [18, 21, 22, 40, 44, 46] and more prominently now Vision Transformers [13, 30] have reached competitive results due to their wider receptive fields. Semantic segmentation requires either global features or contextual interactions to accurately classify, at the pixel level, objects at multiple scales. To this aim, atrous convolutions [3, 5], spatial pyramid pooling modules [61], and networks with attention modules [14, 54, 60] with almost 1 billion parameters were proposed. While these methods enhance the receptive field, achieving good results in terms of performance, they are extremely data hungry and often require strong regularization techniques, such as data augmentation.

2.1. Automatic Data Augmentation

Traditionally, data augmentation requires manual design and domain knowledge. Random cropping, image mirroring, and color distortion are common in natural image datasets like CIFAR-10 and ImageNet. Elastic distortion and affine transformations such as translation and rotation are more common in datasets like MNIST[23] and SVHN [35]. Many methods have been influenced by NAS [63] to find the best dataset-specific set of augmentation policies/strategies [11, 17, 24, 25, 47]. AutoAugment (AA) [11], the first automatic augmentation method, uses reinforcement learning to predict accurate problem-dependant augmentation policies. Despite its success, the approach has the main drawback of running the search on a smaller version of the datasets due to the multi-level search space and the repetitive training, making a strong assumption that the proxy task provides a predictive indication of the larger task. Methods like [12, 27, 34] drastically reduce the parameter space for data augmentation, which allows the methods to be trained on the full dataset. In RandAugmnet (RA) [12] only two hyperparameters are used,

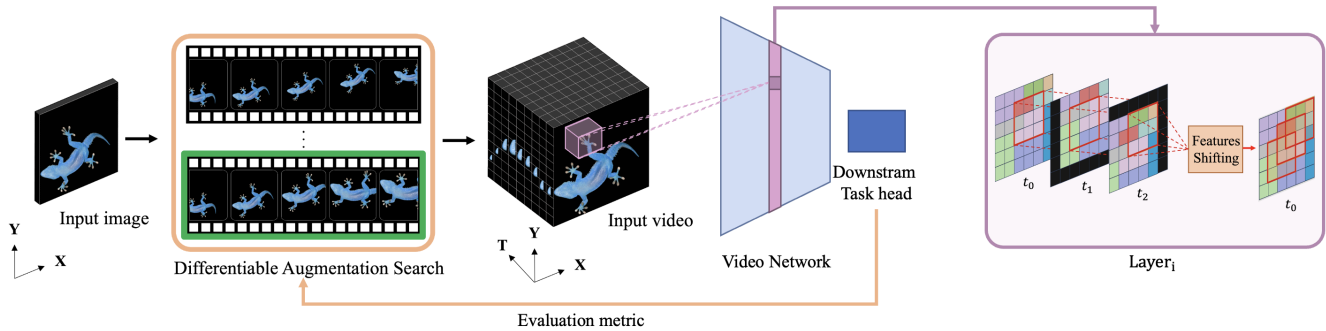


Figure 2. Our method takes an input image and processes it through a DAS cell. The cell, as shown more in detail in Fig. 3, applies all possible transformations defined in the search space and generates an input video. The video is processed through a video network integrated by a temporal shift mechanism, with the goal of shifting the features of adjacent frames. As it can be observed in the pink box, the features shifted and combined as if a kernel $3 \times 3 \times 3$ was applied. As the content derives from transformations of the same image, the result over the original 2D image is a reshaping of the RF. Finally, the predictions for the video input are combined so that the performance for the original 2D task are given back as feedback to the DAS cell.

one controlling the number of augmentations to combine for each image, and the other controlling the magnitude of operations. The downside of RA is that it performs a grid search over a set of augmentation operations incurring up to $\times 80$ overhead over a single training [34]. UniformAugment (UA) [27] and TrivialAugment (TA) [34], instead, propose parameter-free algorithms where hyperparameters are sampled uniformly in the augmentation space. Similarly to UA, we define a continuous augmentations space. However, differently from UA, RA and AA, we do not randomly sample the *ad-hoc* hyperparameters or perform the search on a reduced dataset, but use a continuous search strategy which makes our method fast without the need to reduce drastically the search space or to include user designer bias.

2.2. Enhancing the Receptive Field

In the context of CNNs, the general trend that led to very deep networks is motivated by the seek for a receptive field expansion. The concept of RF is indeed important for understanding and diagnosing how deep CNNs work as it determines what information a single neuron has access to. Over the years, many attempts were made to expand such a field, either by increasing the factors determining its theoretical size, *e.g.* the depth of a neural network and the kernel size, or by providing more context information.

Spatial Domain In the context of image classification and semantic segmentation, using context information from the whole image can significantly help improve the performance. Mostajabi *et al.* [33] demonstrated that by using the “zoom-out” features they can achieve impressive performance for the semantic segmentation task. Liu *et al.* [29] noticed that, although theoretically, the features from the top layers of a neural network should have very large receptive fields, in practice the empirical size is much smaller and consequently not enough to capture the global context. To this aim, they propose to use global averaging to pool the

context features from layers of the neural network, proving empirically that such a method results in a larger empirical RF. In [38], Richter *et al.* question the actual need for a very deep network and propose a method for layer pruning based on the analysis of the size of the RF.

Temporal Domain How to effectively expand the receptive field in the temporal domain has been investigated in video understanding research [26, 32, 41, 42, 51]. 3D CNNs can capture three-dimensional features, however, they are more computationally intensive than 2D CNNs video networks. In [26] the authors implement the Temporal Shift Module (TSM) where a 2D CNN is used as a backbone to extract spatial information. The temporal features are integrated by shifting a fixed amount of channels forward and backward along the temporal dimension. Subsequently, Sudhakaran *et al.* [42] proposed the Gate-Shift-Fuse (GSF) module, a spatiotemporal feature extraction that leverages on a learnable shift of the channels on the temporal dimension as well as channel weighting to fuse the shifted features. TSM and GSF are lightweight and therefore good video network candidates for our image-to-video pipeline in Fig. 2.

In contrast to prior methods that enhance the spatial receptive field by concatenating global or zoomed-out features or that utilize deeper networks, our method extends beyond by aggregating features derived from diverse transformations and by processing them as a video. We employ distinct transformations to generate variations of the input. Given that data augmentation necessitates domain-specific knowledge, we adopt automatic data augmentation.

3. Methods

In this section, we first introduce our new differentiable transformation search process effectively responsible for the reshape of the Receptive Field for image classification

and image segmentation problems (Sec. 3.1). We highlight the benefit of a continuous search space and motivate the importance of a perturbation-based selection technique. We then detail in Sec. 3.2 how we make use of DAS to exploit the power of data, showing that with respect to traditional methods expanding the RF, our effect results in a “reshaping” rather than an expansion. The pipeline of the proposed architecture is shown in Fig. 2.

3.1. Differentiable Augmentation Search

In order to extend images to video and to properly reshape the RF (Fig. 1a), a set of optimal transformations needs to be found. We propose a new general and differentiable approach, later deployed with a restricted search space for our use case, searching for the transformations generating the “best video” to process for a given video network and downstream task. Inspired by Differentiable Neural Architecture Search [28, 53] we define a continuous search space of transformations, which leads to a differentiable learning objective for the joint optimization of the transformations to be applied and the weights of the architecture. Following [28, 64], we search for a computation cell as the initial block that generates the input for the chosen video network. A cell, depicted in Fig. 3, is a directed acyclic graph consisting of an ordered sequence of N nodes. Each node $x^{(i)}$ in the cell is the transformed image, and each directed edge (i, j) is associated with a data augmentation technique. For the deployment of our Differentiable Architecture Search, we defined two search spaces. The first one, similarly to AA and RA methods [11, 12], comprises as set of augmentations Shear X/Y, Translate X/Y, Rotate, AutoContrast, Invert, Equalize, Solarize, Posterize, Color, Brightness, Sharpness, Cutout, and Identity that corresponds to applying no transformation. The cell has two input nodes and a single output node. The size of such a search space is 13^{14} . The second search space, deployed for our purpose, includes Translate X/Y, Scale, Rotate, and Identity. It employs a cell with one single input and output node, having a size of 5^{10} . Given the set \mathcal{T} of candidate transformations, the categorical choice of applying a transformation is relaxed to a Softmax of all possible transformations $t \in \mathcal{T}$. For a given edge (i, j) , the transformation to be applied to the input x is expressed as:

$$\bar{t}^{(i,j)}(x) = \sum_{t \in \mathcal{T}} \frac{\exp(\tau_t^{(i,j)})}{\sum_{t' \in \mathcal{T}} \exp(\tau_{t'}^{(i,j)})} \cdot t(x), \quad (1)$$

where the weights of a transformation are parameterized by a vector $\tau^{(i,j)}$ of dimension $|\mathcal{T}|$. The famous bi-level optimization problem [10] of NAS is re-formularized as:

$$\begin{aligned} \min_{\tau} \quad & \mathcal{L}_{\text{val}}(\tau, \mathbf{w}^*(\tau)) \\ \text{s.t.} \quad & \mathbf{w}^*(\tau) = \arg \min_{\mathbf{w}} \mathcal{L}_{\text{train}}(\mathbf{w}, \tau), \end{aligned} \quad (2)$$

where the model weights \mathbf{w} and the transformations parameters τ are jointly optimized via gradient updates following standard DARTS procedure. We would like to stress that, differently from the original DARTS, here there is no real architecture optimization. The real CNN network that processes the data is fixed and chosen a priori. What we are looking for is a cell where only transformations are applied, and we treat this cell as a sort of “stem” block, placed before the backbone. At the end of the search phase, the best transformations are not chosen by selecting the largest τ value, as this was shown in [53] to be based on a wrong assumption, *i.e.* the τ representing the strength of a transformation agrees with the discretization accuracy at convergence. We rather deploy a perturbation-based approach, where the transformation importance is evaluated in terms of its contribution to the neural network performance. To this aim, for each transformation on a given edge, we mask it while preserving all other transformations, then re-evaluate the cell+CNN. The operation resulting in the greatest reduction in network validation accuracy is identified as the pivotal operation on that edge. Let us indeed consider a cell from a simplified search space composed of only two transformations:

$$\mathbf{I} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{T} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \quad (3)$$

where x_I represents the output of $I \cdot x$, and x_T the translated output $T \cdot x$. Assume m^* to be the optimal feature map, shared across all edges according to the unrolled estimation view. The current estimation of m^* can then be written as:

$$\bar{m}(x) = \frac{e^{\tau_T}}{e^{\tau_T} + e^{\tau_I}} x_T + \frac{e^{\tau_I}}{e^{\tau_T} + e^{\tau_I}} x_I \quad (4)$$

The optimal τ_I^* and τ_T^* minimizing the $\text{var}(\bar{m}(x) - m^*)$ variance between the optimal feature map m^* and the current estimation $\bar{m}(x)$ are:

$$\tau_I^* \propto \text{var}(x_T - m^*) \quad (5)$$

$$\tau_T^* \propto \text{var}(x_I - m^*) \quad (6)$$

We refer to the Supplementary material for a detailed proof. As in the original paper, also for the transformation search space it holds that, from Eq. (5) and Eq. (6), we can see that the relative magnitudes of τ_I and τ_T come down to which one of x_I or x_T is closer to m^* in variance. As x_I comes from the mixed output of a previous edge, and the goal of every edge is to estimate m^* (unrolled estimation) x_I is also directly estimating m^* . x_T is the output of a single transformation instead of the complete mixed output of edge e , so even at convergence it will deviate from m^* . Therefore, if we choose the largest τ as indicating the best transformation, the algorithm will naturally be led to choose identities. Therefore, as on one hand, including the identity in

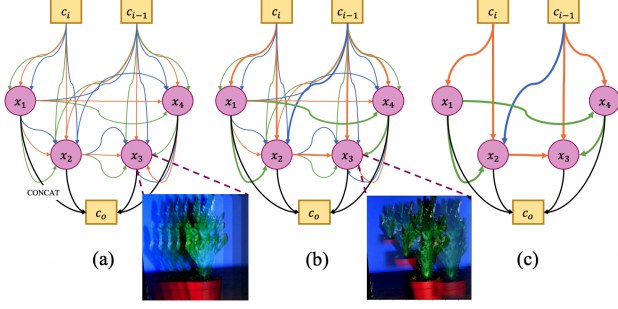


Figure 3. Cell structure in DAS. Multiple operations are defined on each edge, collectively applied to the image, and optimized during training: as the gradients are updated through multiple steps, the τ values associated with each operation change. Fig. 3b depicts one step of such process through thicker edges. In the end, the cell is discretized through a perturbation-based approach, the final operations are chosen and composed (black edges).

the search space is fundamental to remove the a-priori bias that a transformation is always needed, on the other hand, as we show in Sec. 4.2 a cell consisting of only identity transformations leads to poor performance, motivating our transformation selection choice.

3.2. Temporal Data Augmentation

Our goal is to expand and reshape the spatial receptive field. To this aim, we generate videos by picking from a set of transforms, summarized in our search space defined by Translate X/Y, Scale, Rotate, and Identity. We included in our search-space commonly used transformations to model motion in videos [49]. We added scaling, motivated by the general difficulty of segmenting objects at different scale, and identity, to remove the bias that the image *always* benefits from applying a transformation.

Each transformation is applied $\times T$ times, a hyperparameter that defines the length of the generated video. For example, given an image $I(x, y)$, a frame $V(t, x, y)$ of a video applying a vertical translation δy is obtained as:

$$V(t, x, y) = I(x, y - t\delta y) \quad t = 1, \dots, T. \quad (7)$$

As Fig. 2 highlights, the video is given as input to a video network that extracts the features and produces a prediction for the original classification/segmentation task. For the image classification task, we average the predictions for each frame, while for the semantic segmentation task we first “undo” the transformation to preserve the locality concept. Although a general video network can process the video input, as we show in our experiments, to reduce the complexity and keep efficiency in our tasks of interest, we deploy a 2D backbone with a temporal-shift mechanism, *i.e.* GSF, integrated. Such a technique is well-established in the domain of video understanding, allowing to achieve the

performance of 3D CNN but maintaining 2D CNN’s complexity. As the authors of [26, 42] claim, for each inserted GSF, the temporal RF will be enlarged by 2 as if running a convolution with the kernel size of 3 along the temporal dimension. Therefore, part of the features among 3 adjacent frames is mixed. Let us now focus on what the content of those frames is. Fig. 4 gives a proof of concept of our claim for the operations included in our search space. As the content in the adjacent frames is nothing but the same image either translated, scaled, or rotated, the increase in size in the temporal RF can be mapped in an augmentation in size and reshape of the spatial RF. Let us consider the case of a single-path network, for simplicity. The theoretical RF [1] size can be expressed as depending only on the stride s and kernel size k :

$$r_0 = \sum_{l=1}^L ((k_l - 1) \prod_{j=1}^{l-1} s_j) + 1 \quad (8)$$

Let us now consider the situation where one GSF module has been inserted after a 2D convolution with kernel 3×3 , stride $s = 1$. Assuming we are considering the first convolutional layer, the theoretical RF for each frame will have a size of 3. Therefore $r_1^{f_0} = r_1^{f_1} = r_1^{f_2} = 3$. As Fig. 4 shows, the same region now covers different features, that depend on the applied transformation. As these features are mixed through the GSF mechanism, if we map back to the space of frame f_0 , the spatial receptive field will have a size equal to $RF_1 = 3 \times r_1^{f_0} - \bigcap_{i=0}^2 A(r_1^{f_i})$, where $\bigcap_{i=0}^2 A(r_1^{f_i})$ denotes the intersecting area among the receptive field of the different frames. If we define as $L = (x_1, y_1)$ the bottom left corner, and $R = (x_2, y_2)$ the top right corner, the intersection area, for the translation case, between two RFs can be simply found as:

$$A = (\min(x_1, x_1') - \max(y_1, y_1')) \times (\min(x_2, x_2') - \max(y_2, y_2')) \quad (9)$$

where for the case of a translation of t_x along the x-axis and t_y over the y-axis, $x_i' = x + t_x$, $y_i' = y + t_y$. For a rotation of θ degrees counter clockwise, around the point (0,0) (as depicted in 4) $x_i' = x \cos(-\theta) - y \sin(-\theta)$, $y_i' = x \sin(-\theta) + y \cos(-\theta)$. The intersection area calculation however cannot be generalized in this case, as the shape of the intersection polygon depends on the rotation angle. Finally, for the scaling operation the $RF_1 = \gamma^2 \times (r_1^{f_0}) \times \gamma^2 \times (r_1^{f_0})$. Please, note that as some components of the RFs overlap (particularly visible for the case of scaling), what is happening is a reshape of the RF, rather than a re-size, as some parts will contribute more than others. Indeed, we need to recall that there is a difference between the theoretical RF and the *Empirical Receptive Field* ERF, defined as $\frac{\partial y(0,0,0)}{\partial x^0(i,j,z)}$, *i.e.* how much $y(0,0,0)$ changes as $x^0(i,j,k)$ changes by a small amount. Here, (i, j, k) is




Figure 4. RF reshape for translation (top-row) and rotation (bottom-row). In our representation we are assuming that GSF was inserted once, therefore the temporal RF is expanded by 2 and three adjacent frames are considered. The first column shows the features of original image (f_0), while second and third columns show the features of frames obtained applying the transformation. The last column shows the effect of processing the input with a temporal shift mechanism.

the voxel on a pth layer, y is the output and x^0 is the input layer. If we consider the previous example of a 2D convolution with $k = 3 \times 3$, $s = 1$, $t_x = t_y = 1$, $\theta = 30^\circ$, $\gamma = 2$: the new receptive field size for the translation is $RF_1 = 19$, for the rotation is $RF_1 = 14, 19$, and for the scaling is $RF_1 = 144$. As we show in Sec. 4.2, our method reshapes the RF at the cost of a negligible increase # of parameters with respect to standard 2D backbones processing images. However, in terms of memory occupation a trade-off is necessary. We empirically find that expanding the image $\times 5$ times well balances accuracy and memory-occupation. To sum up, we propose a method to expand the spatial receptive field using data augmentation techniques to generate videos. To this aim, we propose a new auto-augmentation method, namely DAS, that, to the best of our knowledge, is the first differentiable auto-augmentation approach proposed in the field of image classification. The “fake videos” are processed by a 2D backbone with an integrated temporal shift mechanism, achieving high performance in video-understanding tasks while keeping 2D CNNs complexity. We finally demonstrate that the increased temporal RF from video processing corresponds to an augmented spatial RF, with the extent determined by the specific transformation used. As we will show in Sec. 4, such a procedure results in lightweight models that reach state-of-the-art and allow reducing the number of parameters typically increased to expand the receptive field, such as kernel size and depth of a network.

4. Experiments

In this section, we investigate the performance of our temporal expansion approach and our automatic augmentation approach for two visual tasks: image classification and semantic segmentation. Best results are highlighted in bold in the tables. We also validate our approach through a series of ablation studies. We refer to the Supplementary materials

Method	# Params	FLOPs	Top-1 Acc.
ResNet-50 [16]	25.6	4.11G	76.30
SE-ResNet-50 [19]	28.1	-	76.90
Inception-v3 [45]	27.2	11.46G	77.12
BnInception [20]	-	-	77.41
Oct-ResNet-50 [6]	25.6	-	77.30
ResNeXT-50 (32×4d) [59]	25.0	8.52G	77.80
Res2Net-50 (14w×8s) [15]	-	-	78.10
ResNet-101 [16]	44.6	7.86G	77.40
ResNet-152 [16]	60.2	11.60G	78.30
SE-ResNet-152 [19]	67.2	-	78.40
ResNeXt-101 (32×4d) [59]	88.8	32.95G	78.80
AttentionNeXt-56 [50]	31.9	-	78.8
ViT-L-32 [13]	304	61.55G	79.66
FFC-ResNet-50 [9]	26.7	5.49G	77.80
FFC-ResNeXt-50 [9]	28.0	5.66G	78.00
FFC-ResNet-101 [9]	46.1	9.23G	78.80
FFC-ResNet-152 [9]	62.6	12.96G	78.90
(Ours) - BnInception	-	-	78.12
(Ours) - Inception-v3	27.3	11.51G	78.66
(Ours) - ResNet-50	25.7	4.2G	79.45
(Ours) - ResNet-101	44.4	7.96G	80.05
(Ours) - ResNet-152	60.4	11.66G	80.13

Table 1. Plugging our method into state-of-the-art networks on ImageNet. The first two sets are top-1 accuracy scores obtained by various state-of-the-art methods, which we transcribe from the corresponding papers. Deeper models are listed in the second set. The third set reports the performances of plugging [9], and last set shows the effect of employing DAS+2D backbone+GSF.




Figure 5. ImageNet results for different Resnet depths.

Method	FLOPs	mIOU
Adelaide_VeryDeep_FCN_VOC [57]	-	79.10
DeepLabv2-CRF [4]	-	79.70
CentraleSupelec Deep G-CRF [2]	-	80.20
HikSeg_COCO [43]	-	81.40
SegModel [39]	-	81.80
TuSimple [52]	-	83.10
Large Kernel Matters [37]	3.7G	83.60
ResNet-38_MS_COCO [58]	12.11G	84.90
PSPNet [61]	16.55G	85.40
DeepLabv3 [3]	49.68G	85.70
(Ours) - ResNet-38_MS_COCO	12.25G	85.70
(Ours) - PSPNet	16.67 G	86.10
(Ours) - DeepLabv3	49.89G	87.00

Table 2. Model comparison with SOTA in PASCAL-VOC-2012.

for all the implementation details, more qualitative results, and for additional experiments on the robustness of DAS.

4.1. Comparison with SOTAs

Image Classification We use ImageNet, Cifar10, Cifar100 and TinyImageNET for the image classification ex-

periments. Tab. 1 shows the comparison of our approach with other state-of-the-art models in ImageNet. The provided results, in the lower part of the table, are obtained running the searching procedure to find the optimal transformation for each independent backbone, *i.e.* BNInception, Inception-v3, ResNet-50-101-152. The best found transformations are given in the Supplementary materials. We employed for this set of experiments *video*-backbones trained from scratch, where we incorporate the best temporal fusion method we found with our ablation studies. We observe an improvement with our approach for every employed 2D backbone using roughly the same parameters, having the strongest boost of 3.15 % in terms of accuracy-wise metric for ResNet-50 architectures. Compared to ViT-L-32 that achieves 79.66 %, we use 1/5 of the parameters and FLOPs. In general, we observe stronger benefits for ResNet-like architectures than Inception-like ones. Moreover, as Fig. 5 highlights, the accuracy of such architectures, with and without our method, when reducing the depth experiences a less steep performance drop. This is reasonable and desired, given the proof of concept provided in Fig. 4. With equal depths, we improve by a fair margin also over FFC [9], a popular method that achieves the expansion of the RF by employing fast fourier convolutions. Similar behaviours are experienced when we compare DAS+temporal expansion with other methods expanding the RF for Cifar10, Cifar100, Tiny and ImageNet datasets, in Tab 4. Of particular notice should be the drop in performance when a bigger kernel is used, probably attributable to a stronger overfitting, and the comparison with the popular dilated convolution (third column), well known for expanding the RF. We improve with respect such a method by a fair margin over each dataset. We attribute this result to the effective *reshaping*, rather than enlargement of the RF.

Image Semantic Segmentation In Tab. 2 and 3 we show the performance of our method applied on Pascal-VOC-2012 and Cityscapes datasets for the image segmentation task. We compare our approach with popular methods employed in semantic segmentation, and observe a considerable gain especially for DAS combined with DeepLab (with a Resnet-101 backbone) when compared with the 2D counterpart. Indeed an improvement of 1.3 % in the mIOU is experienced. The best set of transformations was searched for every different backbone and dataset. We indeed empirically show in Sec. 4.2 how applying not optimized transformations impacts the performance. For Cityscapes, the best results were achieved for ResNeSt backbone, with a 85.1 % mIOU, confirming also a general trend we observed for other datasets. Indeed we experience a much higher improvement when dealing with ResNet architectures than Inception-like modules. This is probably due to the nature of the fusion mechanism we employ, that, when

Method	Backbone	FLOPs	mIOU
FCN [31]	ResNet-101	-	77.02
SETR [62]	ViT-Large	-	78.10
Maskformer [7]	ResNet-101	73G	78.50
NonLocal [55]	ResNet-101	-	79.40
PSPNet [61]	ResNet-101	63G	79.77
Mask2former [8]	ResNet-101	-	80.10
DeepLab-v3 [3]	ResNet-101	92.65G	81.30
DeepLab-v3 [3]	Xception-65	-	82.10
DeepLab-v3 [3]	ResNeSt 101	-	82.90
DeepLab-v3 [3]	ResNeSt 200	285G	83.30
(Ours) PSPNet	ResNet-50	32G	81.21
(Ours) ResNeSt	Xception-65	-	82.40
(Ours) DeepLab-v3	ResNet-101	93G	82.60
(Ours) ResNeSt	ResNet-200	286G	85.10

Table 3. Modal comparison with SOTA in CityScapes.

placed on the skip connections allows better preserving the spatial information. Fig. 6 shows some qualitative results on CityScapes. One can observe a better reconstruction in details when comparing 6c and 6d. This is particularly visible for the reconstruction of people, (see first and third row) and of street lamps. We noticed an over segmentation for PASCAL-VOC-2012 dataset in some cases when applying our methodology. However our method still outperforms all compared state-of-the-art alternatives. We refer the reader to the Supplementary material for some qualitative results on PASCAL dataset and for additional on Cityscapes.

Dataset	Model	Baseline	Dilated	(5 × 5)	Ours
Cifar10	R-18	94.12	94.31	92.16	95.12
	R-50	95.66	96.23	95.51	95.74
	WR-28	96.33	-	-	96.40
Cifar100	R-18	71.20	72.18	70.08	73.10
	R-50	74.82	76.12	75.60	76.71
	WR-28	81.40	-	-	83.06
Tiny	R-18	61.00	62.10	61.03	62.87
	R-50	63.16	64.11	61.12	65.91
	R-101	64.21	65.30	63.60	67.51
ImageNET	R-50	76.30	77.28	76.10	79.45
	R-101	77.40	78.15	76.11	80.05
	R-152	78.30	79.22	78.00	80.13

Table 4. Comparison in terms of accuracy with other methods expanding the RF. “WR” stands for WideResnet.

4.2. Ablation

In this section, we present and discuss ablation studies on Cifar10 and Pascal-VOC-2012. First, we study the effect of each component on the original task accuracy (Tab. 5). We find that the improved accuracy is not a simple matter of augmentation techniques, nor a mere contribution of the fusion mechanism. Indeed, both elements are needed as the simple stack of the same image (third column) does not imply an expansion of the RF, but actually determines a drop of performance probably due to much higher overfit-

	Baseline	Augment	Replica	Ours
Pascal-VOC	85.40	85.51	83.40	86.10
Cifar-10	94.12	94.23	90.11	95.12

Table 5. Ablation on different model components. Augment indicates the usage of additional augmentation derived from our search space *without* temporal fusion. Replica stands for the stack of image copies, that is, video produced with identity transforms and temporal fusion.

	TSM	3D	GSF
Pascal-VOC	85.86 %	77.00 %	86.10 %
	51.32 M	107 M	51.43 M
Cifar-10	94.77 %	92.10 %	95.12 %
	11.18 M	33.17 M	11.20 M
	37.12 M	49 M	37.24 M

Table 6. Ablation on the temporal fusion method given PSPNet backbone for Pascal-VOC and Resnet-18 for Cifar10.

	AA	RA	DAS
Pascal-VOC	82.15	84.00	87.00
Cifar-10	94.35	92.98	95.12

Table 7. Ablation on the auto-augmentation technique given the same searching budget time.

	Random 1	Random 2	DAS
Pascal-VOC	77.25	82.11	86.10
Cifar-10	93.14	94.15	95.12

Table 8. Ablation on the effectiveness of looking for the best transformation. Random 1 is the best transformation found for CityScapes, Random 2 for Cifar100.

ting. Next, in Tab. 6 we ablate on the best temporal fusion mechanism. As expected, we see that TSM and GSF have comparable performance, and number of parameters, with GSF being slightly better. Surprisingly, 3D CNNs perform poorly, maybe because of the small amount of data compared to the needs of 3D CNNs. Next in Tab. 7 we study the impact of the auto-augmentation with our search space. We found that DAS performs best, given a fixed budget searching time of 24 hours decided a priori. We see comparable performance of the three methods for Cifar10 dataset, but we attribute such a behaviour to the relative simplicity of the task itself. Indeed, in Pascal-VOC we experience a much higher difference. Finally, in Tab. 8 we ablate the specificity of the found genotypes for a given dataset, providing quantitative results for the proof of concept provided in Fig. 1a, *i.e.* the effect of random transformation. We check the performance on PASCAL-VOC and Cifar-10 when the genotype to generate the video derives from the search phase for CityScapes (Random1) and for Cifar-100 (Random 2). Interestingly, we see a huge drop in performance when employing a “random genotype”, confirming the need of a carefully designed search procedure.

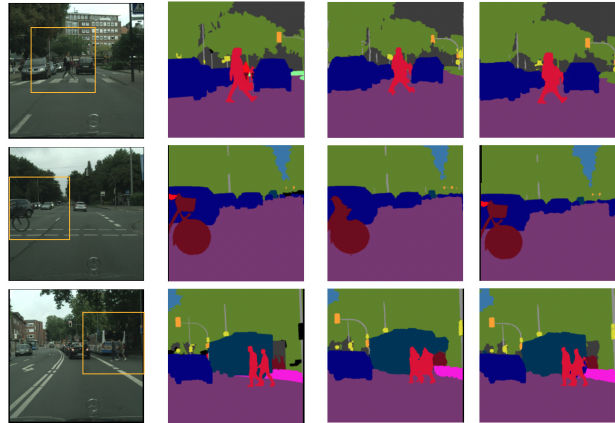


Figure 6. Zoom in of the qualitative results for Cityscapes. The fourth column depicts our results with DAS and DeepLab-v3 integrated with the temporal shift mechanism.

5. Conclusions

We proposed a differentiable auto-augmentation method for the tasks of image classification and semantic segmentation, *i.e.* DAS. We defined a very flexible continuous search space, and employed a perturbation-based selection method to overcome the limitations of previous approaches [11, 12]. We showed that by applying the optimal transformations found by DAS to generate variations of images, we can effectively reshape the RF. This is achieved by processing the new input as videos with a CNN integrated with a temporal shift mechanism that performs feature mixing in time. Our method proposes a new way of handling 2D data to exploit their richness, and investigates how the increase of the receptive field in the temporal dimension impacts the original spatial receptive field. We observed an improvement in terms of accuracy with respect to standard augmentation techniques, for both image classification and segmentation tasks, using different backbones on different datasets. We also successfully reshaped the receptive field, as shown in Fig. 1b, which in terms of qualitative results turned out into more detailed segmentation masks.

Limitations and future work Our method is compact, fast and accurate, but currently limited by its memory footprint. Memory requirement in training grows with number of generated frames, which restricted us to search transformations for 8-frame videos in our experiment. DAS with longer videos might provide further performance boosts. Possible future work include experimenting with other backbone families, such as transformers. Other venue of research could be the adaptation of DAS to video-to-video data augmentation, which would require a proper definition of the search space of transformations. Another interesting direction could be Deep-DAS, where optimal feature map transforms are searched at multiple network layers.

References

- [1] André Araujo, Wade Norris, and Jack Sim. Computing receptive fields of convolutional neural networks. *Distill*, 2019. <https://distill.pub/2019/computing-receptive-fields>. 5
- [2] Siddhartha Chandra and Iasonas Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 402–418. Springer, 2016. 6
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 2, 6, 7
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018. 6
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 2
- [6] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yanis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3435–3444, 2019. 6
- [7] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 7
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 7
- [9] Lu Chi, Borui Jiang, and Yadong Mu. Fast fourier convolution. *Advances in Neural Information Processing Systems*, 33:4479–4488, 2020. 6, 7
- [10] Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of operations research*, 153:235–256, 2007. 4
- [11] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 113–123, 2019. 2, 4, 8
- [12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 2, 4, 8
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 1, 2, 6
- [14] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3146–3154, 2019. 2
- [15] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE transactions on pattern analysis and machine intelligence*, 43(2):652–662, 2019. 6
- [16] 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. 1, 6
- [17] Daniel Ho, Eric Liang, Xi Chen, Ion Stoica, and Pieter Abbeel. Population based augmentation: Efficient learning of augmentation policy schedules. In *International conference on machine learning*, pages 2731–2741. PMLR, 2019. 2
- [18] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2
- [19] 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. 6
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 6
- [21] S Jian, H Kaiming, R Shaoqing, and Z Xiangyu. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision & Pattern Recognition*, pages 770–778, 2016. 2
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 2
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [24] Sungbin Lim, Ildoo Kim, Taesup Kim, Chihon Kim, and Sungwoong Kim. Fast autoaugment. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [25] Chen Lin, Minghao Guo, Chuming Li, Xin Yuan, Wei Wu, Junjie Yan, Dahua Lin, and Wanli Ouyang. Online hyperparameter learning for auto-augmentation strategy. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6579–6588, 2019. 2

- [26] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2019. 3, 5
- [27] Tom Ching LingChen, Ava Khonsari, Amirreza Lashkari, Mina Rafi Nazari, Jaspreet Singh Sambee, and Mario A Nascimento. Uniformaugmt: A search-free probabilistic data augmentation approach. *arXiv preprint arXiv:2003.14348*, 2020. 2, 3
- [28] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2018. 4
- [29] Wei Liu, Andrew Rabinovich, and Alexander C. Berg. Parsenet: Looking wider to see better. abs/1506.04579, 2016. 3
- [30] 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. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, 2021. 2
- [31] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 7
- [32] Chenxu Luo and Alan L Yuille. Grouped spatial-temporal aggregation for efficient action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5512–5521, 2019. 3
- [33] Mohammadreza Mostajabi, Payman Yadollahpour, and Gregory Shakhnarovich. Feedforward semantic segmentation with zoom-out features. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2015. 3
- [34] Samuel G Müller and Frank Hutter. Trivialaugmt: Tuning-free yet state-of-the-art data augmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 774–782, 2021. 2, 3
- [35] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 2
- [36] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1
- [37] Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters—improve semantic segmentation by global convolutional network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4353–4361, 2017. 6
- [38] Mats L. Richter, Schöning J, Wiedenroth A., and Krumnack U. Should you go deeper? optimizing convolutional neural network architectures without training. *arXiv preprint arXiv:2106.12307v2*, 2021. 3
- [39] Falong Shen, Rui Gan, Shuicheng Yan, and Gang Zeng. Semantic segmentation via structured patch prediction, context crf and guidance crf. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1953–1961, 2017. 6
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 2
- [41] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift networks for video action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1102–1111, 2020. 3
- [42] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Gate-shift-fuse for video action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2, 3, 5
- [43] Haiming Sun, Di Xie, and Shiliang Pu. Mixed context networks for semantic segmentation. *arXiv preprint arXiv:1610.05854*, 2016. 6
- [44] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2
- [45] 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. 1, 6
- [46] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 2
- [47] Keyu Tian, Chen Lin, Ming Sun, Luping Zhou, Junjie Yan, and Wanli Ouyang. Improving auto-augment via augmentation-wise weight sharing. *Advances in Neural Information Processing Systems*, 33:19088–19098, 2020. 2
- [48] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 32–42, 2021. 1
- [49] Joost R. van Amersfoort, A. Kannan, Marc’Aurelio Ranzato, Arthur Szlam, Du Tran, and Soumith Chintala. Transformation-based models of video sequences. *ArXiv*, abs/1701.08435, 2017. 2, 5
- [50] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2017. 6
- [51] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pages 20–36. Springer, 2016. 3
- [52] Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding convolution for semantic segmentation. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 1451–1460. Ieee, 2018. 6

- [53] Ruochen Wang, Minhao Cheng, Xiangning Chen, Xiaocheng Tang, and Cho-Jui Hsieh. Rethinking architecture selection in differentiable nas. In *International Conference on Learning Representation*, 2021. 4
- [54] W Wang, J Dai, Z Chen, Z Huang, Z Li, X Zhu, X Hu, T Lu, L Lu, H Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. arxiv 2022. *arXiv preprint arXiv:2211.05778*. 2
- [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. 7
- [56] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022. 1
- [57] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Bridging category-level and instance-level semantic image segmentation. *arXiv preprint arXiv:1605.06885*, 2016. 6
- [58] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019. 6
- [59] 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. 6
- [60] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R. Manmatha, Mu Li, and Alexander Smola. Resnest: Split-attention networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2735–2745, 2022. 1, 2
- [61] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. 2, 6, 7
- [62] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 7
- [63] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 2
- [64] 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. 4