

Dr²Net: Dynamic Reversible Dual-Residual Networks for Memory-Efficient Finetuning

Chen Zhao¹ Shuming Liu¹ Karttikeya Mangalam² Guocheng Qian¹
Fatimah Zohra¹ Abdulmohsen Alghannam¹ Jitendra Malik² Bernard Ghanem¹

¹King Abdullah University of Science and Technology, Saudi Arabia ²UC Berkeley, US

Abstract

Large pretrained models are increasingly crucial in modern computer vision tasks. These models are typically used in downstream tasks by end-to-end finetuning, which is highly memory-intensive for tasks with high-resolution data, e.g., video understanding, small object detection, and point cloud analysis. In this paper, we propose **Dynamic Reversible Dual-Residual Networks**, or **Dr²Net**, a novel family of network architectures that acts as a surrogate network to finetune a pretrained model with substantially reduced memory consumption. Dr²Net contains two types of residual connections, one maintaining the residual structure in the pretrained models, and the other making the network reversible. Due to its reversibility, intermediate activations, which can be reconstructed from output, are cleared from memory during training. We use two coefficients on either type of residual connections respectively, and introduce a dynamic training strategy that seamlessly transitions the pretrained model to a reversible network with much higher numerical precision. We evaluate Dr²Net on various pretrained models and various tasks, and show that it can reach comparable performance to conventional finetuning but with significantly less memory usage. Code will be available at <https://github.com/coolbay/Dr2Net>.

1. Introduction

Large pretrained models play an increasingly crucial role in modern computer vision tasks. These large models, such as ViTs [12] and Swin transformers [30, 31], are pretrained on large-scale datasets [9, 47], by various means such as fully-supervised learning [12, 30], self-supervised learning [14, 20, 36, 41] or vision-language pretraining [38]. They have strong representational capacity due to the large model scale and the data scale, and therefore become indispensable for various downstream tasks [1, 4, 18, 37, 40, 43, 46].

Although these pretrained large models have shown

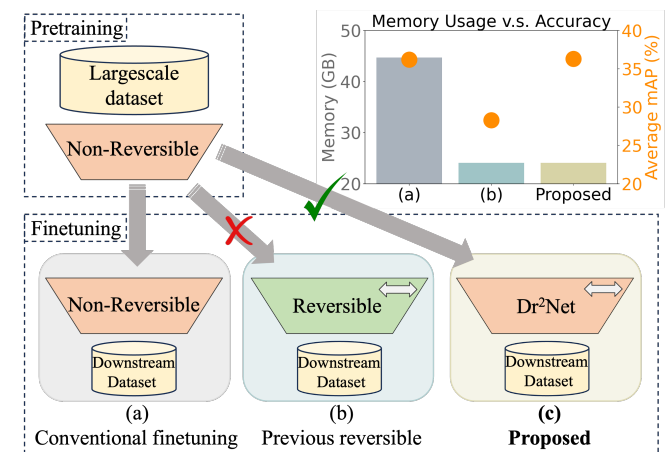


Figure 1. **Comparison of different ways of finetuning from pretrained non-reversible models.** (a) *Conventional finetuning* uses the same non-reversible architecture in the downstream task, initialized with the pretrained parameters. It consumes high GPU memory. (b) *Previous reversible* methods (e.g., [16, 33, 46]) cannot finetune from pretrained non-reversible models on the downstream task due to architecture discrepancy. They show reduced accuracy when training from scratch on the downstream. (c) *Our proposed Dr²Net* can directly finetune from pretrained non-reversible networks, significantly saving memory while preserving accuracy. The top-right chart illustrates memory usage and accuracy for temporal action detection on ActivityNet-v1.3 [5] using VSGN [44] and Video Swin [31].

good generality, they need to be end-to-end finetuned on specific downstream tasks to reach an optimal performance [4, 8, 27, 28, 37, 43, 46]. End-to-end finetuning refers to training the backbone, which is initialized with a pretrained model, simultaneously with the task-specific network during finetuning. For example, it is a common practice in image object detection that the backbone is initialized from a model pretrained on ImageNet classification [9] and finetuned end-to-end on the object detection datasets [43]. For the task of video temporal action localization, recent research has shown performance boost of using end-to-end finetuning compared to frozen-backbone

finetuning [8, 27, 28, 46]. For self-supervised pretrained models such as MAE [20], end-to-end finetuning is required to reach even a decent performance for downstream tasks.

However, end-to-end finetuning is memory intensive, especially for those large models on a task with high-dimension or high-resolution data, as shown in Fig. 1 (a). For example, in long-form video understanding tasks, *e.g.*, temporal action localization [8, 27, 46], thousands of video frames need to be processed at a time for long-term reasoning. Without dramatically downscaling the resolution, it is even impossible to finetune a Video Swin - large model with a video of 30 seconds in the largest GPU, *i.e.*, A100 with 80 GB memory [46]. Therefore, reducing GPU memory consumption is a vital problem in finetuning large models.

Recently, reversible networks have demonstrated their efficacy in significantly reducing memory consumption during training [16, 22, 24, 33, 46]. They can reconstruct intermediate activations from network output, and therefore don't need to store those activations in memory during the forward process. However, existing reversible networks [16, 22, 24, 33] are not able to leverage pretrained models, and have to be trained from scratch, which leads to inferior performance as shown in Fig. 1 (b). While the more recent work Re²TAL [46] proposed a rewiring strategy enabling the reuse of pretrained model architectures and parameters, it still requires pretraining the reversible model before finetuning it on the downstream task to maintain performance. A major challenge in directly fine-tuning reversible networks from pretrained models is the inherent architectural disparity. The majority of existing pretrained models are designed as non-reversible networks, making direct transfer learning to the distinctly different architecture of reversible networks in the downstream tasks challenging.

To reduce memory consumption without compromising performance when finetuning pretrained non-reversible models on downstream tasks, in this paper, we propose a family of network architectures, dubbed as **Dynamic Reversible Dual-Residual Networks** or Dr²Net. Dr²Net acts as a surrogate backbone network during finetuning, and can be seamlessly initialized from pretrained non-reversible models. Dr²Net is essentially a super network encompassing both the pretrained non-reversible architecture and the downstream reversible architecture. It employs two types of residual connections: one preserves the residual structure of the pretrained non-reversible architecture, while the other facilitates reversibility. By applying two distinct coefficients to these residual connections, we can control the network's proximity to either architecture. During finetuning, we dynamically update the coefficients such that the network seamlessly transitions from the pretrained non-reversible model to a reversible network of increased numerical precision (referred to as a robust reversible network). This design effectively bridges the architectural gap

between the two types of networks.

We summarize our contributions as follows.

- We propose a novel family of network architectures dubbed as Dynamic Reversible Dual-Residual Networks (Dr²Net) to finetune any pretrained model with substantially reduced memory consumption.
- We introduce a dynamic finetuning strategy to seamlessly transition any network to a robust reversible network, achieving performance comparable to conventional finetuning while conserving memory.
- We have shown the effectiveness of Dr²Net on various pretrained models such as Swin [30] and ViT [12], and a broad range of vision tasks such as temporal action detection [42]. Dr²Net significantly reduces memory while preserving accuracy.

2. Related Works

2.1. Large pretrained models

Large pretrained models [20, 31, 36, 41], due to their large model scales and their large-scale training data, have demonstrated impressive performance in various computer vision tasks. Different pretraining mechanisms have been explored in the literature. Fully-supervised classification, *e.g.*, image classification [12, 30] on ImageNet [9] and video action classification [13, 31] on Kinetics [47], is a common pretraining task when the data categories are available. When annotation is scarce, self-supervised learning, *e.g.*, MAE [20], VideoMAE [41] and DINOv2 [36], is an effective way to leverage large-scale unlabeled data. These models can scale up more easily by utilizing the vast amount of images and videos out there without human annotation. If paired language descriptions for the vision data are available, vision-language pretraining can be utilized, *e.g.*, CLIP [38] and Frozen [3].

All these types of pretrained models can benefit downstream tasks through finetuning. It has been shown that finetuning from large pretrained models achieves significantly improved performance than training from scratch for various downstream tasks [4, 37, 41, 43, 46]. However, the majority of existing pretrained models are non-reversible, and consume a large amount of GPU memory when used for downstream finetuning. In this paper, we propose a novel family of reversible networks for downstream finetuning, which can directly leverage these pretrained models.

2.2. Memory-efficient training

The computational demands, *e.g.* GPU memory, impede deep neural network training. Various techniques have been proposed to mitigate this issue. For instance, mixed precision training [34] reduces the numerical precision of certain model layers while maintaining performance, thereby lowering memory usage. Another approach, activation check-

pointing [7], stores only specific intermediate activations in the forward pass and recomputes the others during back-propagation. However, their memory costs scale linearly or sublinearly with the number of network layers, posing challenges for deeper networks. Besides these architecture-agnostic approaches, efforts to design specific memory-efficient networks have also been fruitful. An exemplary case is the reversible network [16, 32], which requires storing only the final feature map during forward propagation. This storage requirement remains constant irrespective of network depth. In backpropagation, the reversible network efficiently reconstructs intermediate feature maps from deep to shallow layers, offering a more scalable solution than activation checkpointing. In this paper, we adopt the idea of reversible networks for memory efficiency.

2.3. Reversible networks

Reversible networks originated from the idea of invertible transformations in NICE [10, 11], which inspires subsequent architectures proposed for various purposes, for example normalizing-flow based image generation [21, 23], signal reconstruction [29, 35] and memory efficiency [16, 32, 46]. RevNet [16] adapts the NICE transformation for ResNets [19] and proposes a reversible backpropagation algorithm that significantly reduces the GPU training memory cost. RevViT [32] further adapts the NICE transformers to Vision Transformers [12] and achieves performance parity across a variety of tasks. Re²TAL [46] proposes a method to rewire a pretrained non-reversible backbone into a reversible backbone. But their proposed method still needs fine-tuning the reversible network on the pretraining task using the pre-training dataset. However, quite often, downstream practitioners do not have ready access to the pre-training dataset or the pre-training implementation and recipes. Further, it fails on finetuning from self-supervised learned models, such as VideoMAE [41]. Reversible networks are an effective approach for conserving memory, yet existing methods are not able to transfer the parameters from a non-reversible network to a reversible network. In this paper, we propose a new type of reversible networks, which enable directly finetuning from parameters of pretrained non-reversible networks in the downstream tasks.

2.4. Memory-intensive tasks

Many computer vision tasks, which involve high-dimension or high-resolution data, are highly memory intensive, such as long-form video understanding, small object detection.

Long-form video understanding. Temporal action detection (TAD) [8, 25, 28, 39, 42, 45, 46] is a typical long-form video understanding task. It requires reasoning among a large number of video frames, and therefore uses a lot of GPU memory. We cannot even feed a video of 30 seconds into the largest GPU without significantly downscaling the

video resolution. To enable training with a long sequence of video frames, most methods in the literature adopt the feature-based mechanism, where they freeze the gigantic pretrained backbone and only train the TAD-specific layers [25, 42, 44]. However, this inevitably sacrifices accuracy. Some recent methods propose to do end-to-end training by sampling a subset of the video snippets for processing (e.g., TallFormer [8]) or for back-propagation (e.g., ETAD [27]). However, these snippet-sampling based methods require each video snippet to be independently encoded, and cannot perform global temporal aggregation.

Object detection in large images. To accurately detect small objects, state-of-the-art object detectors (e.g., DINO [43]) rely on a large-resolution input image, such as 1024×1024 , requiring massive GPU memory for training. Consequently, limited model sizes or batch sizes can be utilized, restricting the detection accuracy. Additionally, object detection heavily relies on the pretrained image backbone, which is utilized to initialize the detection backbone and then finetuned for the detection task. Studies show that finetuning from models pretrained on larger image classification datasets [20] or detection datasets [30, 43] can significantly boost the detection performance.

In this work, we use our proposed Dr²Net to dramatically reduce memory consumption of end-to-end finetuning for these memory-intensive tasks. Using the saved memory, higher input resolutions or larger models can be utilized to reach higher performance.

3. Methodology

3.1. Problem formulation

Given a pretrained model, which is usually a non-reversible neural network, e.g., Video Swin [31] trained on Kinetics [47], we denote its architecture as \mathcal{M}_n (Fig. 2 (a)), and its parameters as θ_n . Our objective is to finetune the model on a downstream task, such as video temporal action detection [8, 27, 46], in a memory-efficient manner. Typically, conventional finetuning involves transferring both the architecture and the parameters from the pretrained model to the downstream task, as illustrated in Fig. 1 (a). Concretely, the same architecture \mathcal{M}_n is directly utilized as the backbone architecture in the downstream task, and it is initialized with the parameter values θ_n during the finetuning process.

However, finetuning on a downstream task with high data dimension or resolution is memory-intensive. To mitigate this, instead of using the same architecture \mathcal{M}_n in the downstream task, we propose to transform the pretraining architecture \mathcal{M}_n to a reversible one \mathcal{M}_r as the downstream backbone, and initialize its parameters with the same values θ_n for finetuning. This approach raises two key questions: (1) How do we transform the architecture to a reversible one that can seamlessly reuse the parameter values θ_n ? (2)

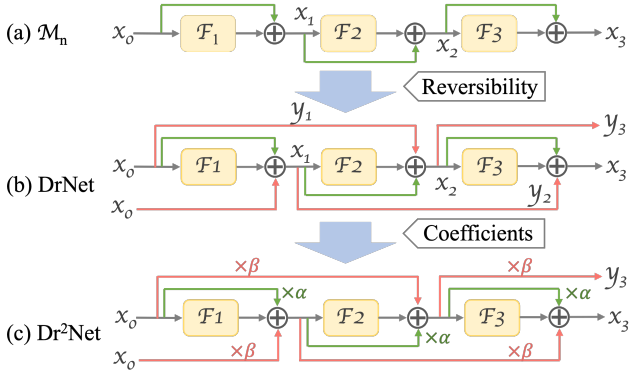


Figure 2. **Transforming a pretrained non-reversible network architecture \mathcal{M}_n into our proposed Dr^2Net .** (a) \mathcal{M}_n : the pretrained non-reversible network with three blocks $\mathcal{F}_i, i = 1, 2, 3$. Considering that most contemporary networks have residual connections, we illustrate the network with residual connections in the figure (green arrows), though our method doesn't restrict \mathcal{M}_n to be residual networks. (b) **DrNet**: a reversible network obtained by adding a new group of residual connections (pink arrows) to \mathcal{M}_n . (c) **Dr²Net**: our proposed reversible network obtained by adding coefficients α and β to the two groups of residual connections respectively. Dr^2Net is equivalent to \mathcal{M}_n when $\alpha = 1$ and $\beta = 0$. Note that the blocks \mathcal{F}_i can be of any architectures following [46], and there can be any number of \mathcal{F}_i blocks in each network.

How do we effectively finetune the reversible network in a memory-efficient setting? Sec. 3.2 and Sec. 3.3 will address these two questions respectively.

3.2. Dynamic Reversible Dual-Residual Networks

To transform a network into a reversible one, recently, Re^2TAL [46] proposed to rewire the residual connections of a non-reversible network (Fig. 2 (a)) to obtain a reversible architecture. During its rewiring process, though the basic blocks \mathcal{F}_i are maintained, the macro architecture has a significant change. Consequently, the obtained reversible network is mathematically a different function from the original network, and cannot be directly finetuned from the parameters θ_n of the original network on the downstream task.

What if we can maintain the macro architecture during the network transformation? Actually, initializing network parameters from a different network has been studied in the literature [3, 6]. In I3D [6] and Frozen [3], researchers attempted to initialize video networks from parameters of pretrained image networks. They adopted a parameter initialization mechanism that makes the video network equivalent to the image network before finetuning by initializing the extra parameters in the video network to certain values. Inspired by these, we propose to transform the architecture with minimum modification, and ensure equivalency between the pretrained network and the downstream network at the beginning of finetuning.

To obtain a reversible downstream network, instead of rewiring the residual connections in the original network, we add new residual connections, as illustrated in Fig. 2 (b). Note that since most contemporary networks have residual connections, we illustrate the pretrained architecture \mathcal{M}_n as a residual network following [46], though our method does not restrict \mathcal{M}_n to residual networks. At each block \mathcal{F}_i , we add a new residual connection, the pink arrow in Fig. 2 (b), that skips two blocks \mathcal{F}_i and \mathcal{F}_{i+1} . We replicate the original input x_0 as the input to the first new residual connection to form a second pathway. The obtained network is a reversible **Dual-Residual Network**, **DrNet** for short, whose reversibility will be detailed later in this section and proved in the appendix. **DrNet** preserves the original residual connections. However, it still has an architectural discrepancy from the original network \mathcal{M}_n due to the newly introduced residual connections (pink arrows).

To enable initializing the reversible network to be the equivalent of the pretrained \mathcal{M}_n , we introduce two coefficients on the two types of residual connections respectively. We use $\alpha \in [0, 1]$ on the original residual connections (green), and $\beta \in [0, 1]$ on our newly added residual connections (pink). With these two coefficients, we actually obtain a family of reversible networks with α and β set to different values. These two coefficients can be dynamically adjusted during finetuning (see Sec. 3.3), so we call the obtained new architecture **Dynamic Reversible Dual-Residual Networks**, **Dr²Net** for short, as illustrated in Fig. 2 (c).

We use Dr^2Net for downstream finetuning with the parameters θ_n of the pretrained network \mathcal{M}_n as initialization. When initializing Dr^2Net , we make $\alpha = 1$ and $\beta = 0$, such that it becomes exactly the same architecture as the pretrained network \mathcal{M}_n . In this way, we can seamlessly initialize Dr^2Net using the parameters θ_n of the pretrained network \mathcal{M}_n . When we make $\alpha = 1$ and $\beta = 1$, Dr^2Net becomes **DrNet** above; when we make $\alpha = 0$ and $\beta = 1$, Dr^2Net becomes an architecture as in Re^2TAL [46].

We mathematically formulate the computation of the i^{th} module in Dr^2Net as follows

$$\begin{cases} y_i = \beta \times x_{i-1} \\ x_i = \mathcal{G}_i(x_{i-1}) + y_{i-1}, \end{cases} \quad (1)$$

where $y_0 = x_0$, and $\mathcal{G}_i(x_{i-1}) = \mathcal{F}_i(x_{i-1}) + \alpha \times x_{i-1}$. If the pretrained network \mathcal{M}_n doesn't have residual connections, $\alpha = 0$. We can observe that this Dr^2Net module is reversible as long as $\beta \neq 0$. Its reverse computation is formulated as follows

$$\begin{cases} x_{i-1} = y_i / \beta \\ y_{i-1} = x_i - \mathcal{G}_i(x_{i-1}). \end{cases} \quad (2)$$

Recall that we need to make $\beta = 0$ to make Dr^2Net equivalent to the pretrained network \mathcal{M}_n , which conflicts the $\beta \neq 0$ requirement here. We will discuss this in Sec. 3.3.

Due to the reversibility of each module, all the intermediate activations x_i and y_i can be reconstructed from the output, and hence don't need to be cached in memory [16, 33, 46].

3.3. Memory-efficient finetuning

To finetune our reversible network Dr²Net in a memory-efficient manner, we clear intermediate activations from memory during the forward process, and customize the back propagation to reconstruct those activations from the output using Eq. 2 following the implementations in [16, 33]. As [16] pointed out, in reversible networks, though the activations can be exactly reconstructed when done in exact arithmetic, numerical error may be accumulated during back propagation due to floating point computation with limited precision. If the numerical error is within a certain level, it does not affect the performance; otherwise, training will be impaired. In this subsection, we will discuss how to effectively finetune our Dr²Net in a memory-efficient manner with minimum influence from the numerical error.

3.3.1 Vanilla finetuning

As described in Sec 3.2, we set $\alpha = 1$ and $\beta = 0$ to make Dr²Net equivalent to the pretrained network \mathcal{M}_n at the beginning of the finetuning. However, β can not be 0 because it will be used in the denominator in the reverse process, as shown in Eq. 2. To circumvent this constraint, we use a small value for β instead. During our experiments, we find that when β is too small, the numerical error will corrupt the training. As a tradeoff between numerical precision and the resemblance of Dr²Net and \mathcal{M}_n , we make $\beta = 0.1$.

Can we keep using the same architecture of Dr²Net with $\alpha = 1$ and $\beta = 0.1$ throughout the finetuning process? To answer this question, we need to know whether and how the values of α and β will influence the numerical error of Dr²Net finetuning. To this end, we carry out a study on the relationship between the back propagation error and the values of α and β . To train Dr²Net, besides the memory-efficient training that clears intermediate activations and reconstructs them during back propagation, we can also store its intermediate activations in GPU memory and perform common back propagation, which doesn't save memory but provides accurate gradient values not impacted by numerical errors as a reference. We compare the gradient values between the two ways of training with different α and β values. In Fig. 3, we plot the gradient error levels (*i.e.* magnitudes) of our Dr²Net with Video Swin-tiny [31] under FP32.

We can see that the error level is the lowest 10^{-12} when $\alpha = 0$ and $\beta = 1$ on the bottom-left corner; it is the highest 10^{-1} when both α and β are close to 1 at the same time on the bottom-right corner. The error level in the middle area is around 10^{-7} , which is already acceptable. Our selection of $\alpha = 1$ and $\beta = 0.1$ on the top-right corner to initialize the




Figure 3. **Gradient error levels with different α and β values.** The scales on the right colorbar represent $10^{-12} \sim 10^{-2}$. When $\alpha = 1$ and $\beta = 0.1$ (top-right) at the beginning of finetuning, the error level is 10^{-5} . The errors are the smallest when $\alpha = 0$ and $\beta = 1$ (bottom-left) and the largest when $\alpha = 1$ and $\beta = 1$ (bottom-right). The error level in the middle area 10^{-7} is already acceptable. The blue arrows represent an ideal evolution path of the two coefficients over the finetuning process: progressively approaching the values that produce acceptable gradient errors.

architecture has an error level of 10^{-5} , which is not detrimental to the training, but doesn't give precise results. That indicates that if we keep using these values, it is probably very hard to reach the optimal solution due to the imprecise gradients (see Tab. 5, 6, and 7 in Sec. 4.2). Therefore, we need a *dynamic finetuning* mechanism to adjust the coefficient values to a low-error point during finetuning.

3.3.2 Dynamic finetuning

As mentioned above, we identify two key factors that impact the effectiveness of Dr²Net finetuning: (1) the model's proximity to the pretrained network \mathcal{M}_n , and (2) gradient precision of the customized back-propagation. The former needs a high α value and a low β value (the top-right corner in Fig. 3), whereas the later needs a low α value and a high β value (the bottom-left corner in Fig. 3). This presents an apparent contradiction in the finetuning requirements.

Upon further examination, we find that the relative importance of these factors changes over the course of finetuning. Initially, Factor (1) is critical as Dr²Net begins finetuning with the pretrained parameters θ_n . However, its significance diminishes as the network adapts over successive finetuning epochs. Conversely, Factor (2) becomes crucial in the later stages when seeking a precise solution, though it is less critical during the initial, more chaotic phases.

First, at the beginning of finetuning, $\lambda \rightarrow 0, \alpha \rightarrow 1$, *i.e.*, the architecture needs to be as close as possible to the original pretraining network \mathcal{F} , to have a matched initialization. But this becomes less important at later finetuning

Table 1. **Downstream tasks experimented in this paper.** These tasks involve various pretraining methods, including fully-supervised classification and self-supervised learning methods such as MAE [20]. They also use different backbones, e.g., Video Swin [31], ViT [12].

Data type	Pretraining	Backbone	Downstream task	Downstream method	Downstream dataset
Video	Classification	Video Swin - tiny [31]	Temporal action detection	VSGN [44]	ActivityNet-v1.3 [5]
	Classification	Video Swin - tiny [31]	Refer. video object segment.	MTTR [4]	A2D-Sentences [15]
	VideoMAE [41]	Video ViT - small [41]	Action recognition	VideoMAE [41]	SthSth-v2 [17]
Point cloud	MAE [20]	ViT - small [12]	Point cloud segmentation	Pix4Point [37]	S3DIS [2]
Image	Classification	Swin - tiny [30]	Object detection	DINO [43]	MS-COCO [26]

stages when the network has already evolved over iterations of training. **Second**, at later stages when the optimal solution is being sought, $\lambda \rightarrow 1, \alpha \rightarrow 0$, i.e., the architecture needs to be as close as possible to the reversible network, to ensure precise back propagation. But this doesn't matter as much at early stages when the loss itself is high.

Based on these analyses, we propose a dynamic finetuning mechanism for Dr²Net. At the beginning of finetuning, we use $\alpha = 1, \beta = 0.1$, and initialize Dr²Net with the parameters θ_n of the pretrained model \mathcal{M}_n . Then during the finetuning process, we progressively decrease the value of α and increase the value of β until an acceptable gradient error level is reached, a point we call an update end point, e.g. 10^{-7} as shown in Fig. 3. After that, we use the fixed α and β values for the rest of the training iterations. The next question is how to schedule the decrease of α and the increase of β to have the lowest impact to the accuracy.

Updating schedule: α first or β first? There are three options of updating the values of α and β : (1) fully updating α before updating β , (2) fully updating β before updating α , and (3) simultaneously updating α and β . Seen from Fig. 3, the shortest path connecting the start point (the top-right corner) and the update end point is along the diagonal signified by the blue arrows. This path reaches the update end point, which has a desired gradient error level, at earlier iterations than the paths in vertical or horizontal directions. This diagonal path corresponds to the third option — simultaneously updating α and β . We will compare the performance of the three options in Sec. 4.2.

Updating policy: what functions to compute the values of α and β ? We consider α and β as functions of training iterations, α being a monotonically non-increasing function and β being a monotonically non-decreasing function. α always starts from 1 and β always from 0.1. They end changing at the update end point, e.g. $\alpha = 0.3, \beta = 0.7$ as illustrated in Fig. 3. During our experiments, we find that simple linear functions are effective (see Sec. 4.2), and therefore employ linear functions for all the tasks.

Updating frequency and end epoch. Under our updating policy, there are two hyper-parameters. The first one is the updating frequency η , which means that the two coefficients are updated every η epochs / iterations. The second one is the end epoch τ , which means that the two coeffi-

cients reach the update end point at the τ^{th} epoch. We provide the comparison of different η values for the task of action recognition with VideoMAE pretrained model [41] in Sec. 4.2, and have the choices of the two hyper-parameters for all tasks in the appendix.

4. Experiments

We conducted extensive experiments on various tasks to show the effectiveness of our proposed Dr²Net, and summarize them in Tab. 1. We target 5 different kinds of vision tasks that require high-dimensional data such as videos, or high-resolitional images such as point cloud. As listed in the table, these tasks use different downstream datasets, and adopt different backbones, including Swin [30], ViT [12], Video Swin [31] and Video ViT [41], which have been pretrained in different ways, such as fully supervised classification and self-supervised learning MAE [20]. We provide the implementation details of these tasks in the appendix.

4.1. Effectiveness of Dr²Net

In Tab. 2 and Tab. 3, we show the effectiveness of our Dr²Net in memory saving as well as in accuracy preservation, by comparing to conventional finetuning and other reversible methods on all the observed tasks listed in Tab. 1. Conventional finetuning uses the same non-reversible network as pretraining, and therefore consumes a large amount of GPU memory with end-to-end finetuning (Row 2). If we freeze the backbone and only train the downstream task-specific layers (Row 1), memory usage is dramatically reduced, but at the same time accuracy is also significantly impaired. Previous reversible models (e.g., [46]) cannot directly finetune from pretrained non-reversible model, and have to train from scratch (Row 3). Re²TAL [46] supports reusing the architecture of the pretrained model, and with it we tried *hard finetune* by initializing the rewired reversible model using the pretrained parameters (Row 4). We can see that training the reversible model from scratch on the downstream tasks lead to obviously inferior performance. With *hard finetune*, better performance is achieved, but there is still a big gap from conventional end-to-end finetuning.

Seen from Tab. 2 and Tab. 3, our proposed Dr²Net (Row 5), saves 46.1%, 56.6%, and 79.5% memory for the three video tasks respectively, and saves 30.6% and 44.4% mem-

Table 2. **Memory and accuracy comparison on different video understanding tasks.** Conventional: conventional non-reversible backbone; Reversible: previous reversible backbone [46]; Hard finetune: directly initializing the reversible backbone using pretrained parameters. “mAP”: mean average precision, “Acc”: accuracy.

Downstream training		Temporal action detection [44]		Refer. video object segmentation [4]		Action recognition [41]	
		Avg. mAP	Memory	mAP	Memory	Top-5 acc	Memory
Conventional	Frozen backbone	34.4%	12.2GB	41.2%	9.1GB	29.9%	2.4GB
	End-to-end	36.2%	44.7GB	44.1%	41.5GB	90.3%	29.3GB
Reversible [46]	From scratch	28.3%	24.1GB	/	18.0GB	56.8%	6.0GB
	Hard finetune	35.4%	24.1GB	42.1%	18.0GB	78.9%	6.0GB
Dr²Net	End-to-end	36.3%	24.1GB	42.9%	18.0GB	89.0%	6.0GB

* We reproduced the results of conventional end-to-end finetuning, using the official code of each method for fair comparison.

Table 3. **Memory and accuracy comparison on point cloud segmentation and object detection.** Convent.: conventional non-reversible backbone; Reversible: previous reversible backbone [46]; Hard finetune: directly initializing the reversible backbone using pretrained parameters. “mIoU”: mean intersection over union, “mAP”: mean average precision. Note that *frozen backbone* for *conventional* doesn’t save much memory for point cloud segmentation since the downstream method Pix4Point [37] has added parameters before the network backbone.

Downstream training		Point cloud seg. [37]		Object dect. [43]	
		mIoU	Memory	mAP	Memory
Convent.	Frozen backb.	62.0%	21.2GB	49.7%	26.9GB
	End-to-end	69.6%	22.5GB	51.3%	54.0GB
Reversible	From scratch	65.7%	15.6GB	38.7%	30.0GB
	Hard finetune	62.5%	15.6GB	49.6%	30.0GB
Dr²Net	End-to-end	68.1%	15.6GB	51.3%	30.0GB

* We reproduced the results of conventional end-to-end finetuning, using the official code of each method for fair comparison.

Table 4. **Training time comparison of Dr²Net to conventional end-to-end finetuning.** The numbers are training time per epoch.

Task	TAD [44]	RVOS [4]	AR [41]	PCS [37]	OD [43]
Conventional	198 min	24 min	50 min	178 sec	206 min
Dr ² Net	261 min	29 min	89 min	186 sec	216 min

ory for point cloud segmentation and object detection respectively. Across these experiments, our Dr²Net reaches comparable accuracy to the original network while significantly reducing memory consumption when finetuning end-to-end. In these experiments, we use the smallest network variants for each type of backbone due to limited computational resources. Note that the memory saving will be more significant with deeper networks since reversible networks use constant memory regardless of network depths [33, 46].

Theoretically, the reversible training adds about 33% more operations, as pointed out in RevNet [16], but the actual latency can be smaller, varying among tasks. Table 4 compares the training time of conventional end-to-end training and Dr²Net for all the five tasks.

4.2. Ablation Study and Design Analysis

We perform the following ablation study and design analysis on multiple tasks to validate our design choices.

Table 5. **Ablation study of dynamic finetuning on temporal action detection [46].** Vanilla finetuning uses fixed values of $\alpha = 1$ and $\beta = 0.1$ throughout finetuning, our Dr²Net uses dynamic finetuning, which updates the two coefficients dynamically.

Method	0.5	0.75	0.95	Avg. mAP
Vanilla finetune	52.25%	35.86%	10.01%	35.44%
Dynamic finetune (Dr²Net)	53.24%	36.97%	10.16%	36.27%

Dynamic finetuning ablation is shown in Tab. 5, Tab. 6, and Tab. 7 for the tasks of temporal action detection [44], action recognition [41], and point cloud segmentation [37], respectively. Our Dr²Net uses dynamic finetuning that updates the values of the two coefficients α and β during the finetuning process, as described in Sec. 3.3. We compare it to using vanilla finetuning, which uses fixed values of $\alpha = 1$ and $\beta = 0.1$ throughout the finetuning process. From the tables, we can see that Dynamic finetuning leads to obviously higher accuracy than vanilla finetuning for all the tasks. The advantage of dynamic finetuning is significantly evident with the VideoMAE [41] pretrained model, as shown in Tab. 6. Without dynamic finetuning, action recognition using the VideoMAE pretrained model (Row 1 in Tab. 6) is even worse than training from scratch (Row 3 in Tab. 2).

α and β **updating schedules** are compared in Tab. 8, Tab. 9, and Tab. 10 for the tasks of action recognition [41], object detection [43], and point cloud segmentation [37], respectively. Our Dr²Net updates α and β simultaneously instead of finishing updating one before updating the other, as described in Sec. 3.3. We compare our simultaneous updating schedule to the following two schedules with the same updating frequency: (1) update α first until it reaches the updating end point, and then update β (Row 1); (2) update β first until it reaches the updating end point, and then update α (Row 2). Empirically, we find that the gradient error level of 10^{-7} (3) can preserve the accuracy very well, therefore we consider the α and β values at this point as the update end point. From the three tables, we can tell that our simultaneous updating schedule gives higher performance than the other schedules.

Updating frequency η is studied in Tab. 11 for the task

Table 6. **Ablation study of dynamic finetuning on action recognition with VideoMAE pretrained model [41].** Vanilla finetuning uses fixed values of $\alpha = 1$ and $\beta = 0.1$ throughout the finetuning, our Dr²Net uses dynamic finetuning, which updates the two coefficients dynamically.

Method	Top-1 Acc	Top-5 Acc
Vanilla finetune	26.48%	53.08%
Dynamic finetune (Dr²Net)	64.57%	89.01%

Table 7. **Ablation study of dynamic finetuning on point cloud segmentation [37].** Vanilla finetuning uses fixed values of $\alpha = 1$ and $\beta = 0.1$ throughout the finetuning, our Dr²Net uses dynamic finetuning, which updates the two coefficients dynamically.

Method	mIoU
Vanilla finetune	57.57%
Dynamic finetune (Dr²Net)	68.13%

Table 8. **Comparison of α or β updating schedules on action recognition with the VideoMAE pretrained model [41].** ‘‘Acc’’ means accuracy. Simultaneously updating α and β leads to the highest accuracy.

Schedule	Top-1 Acc	Top-5 Acc
α first, β second	60.51%	86.64%
β first, α second	58.40%	85.30%
α and β simultaneously (Dr²Net)	64.57%	89.01%

Table 9. **Comparison of α or β updating schedules on object detection [43].** Simultaneously updating α and β leads to the highest accuracy.

Schedule	mAP
α first, β second	50.2%
β first, α second	50.7%
α and β simultaneously (Dr²Net)	51.3%

Table 10. **Comparison of α and β updating schedules on 3D point cloud segmentation [37].** Simultaneously updating α and β leads to the highest accuracy.

Schedule	mIoU
α first, β second	66.40%
β first, α second	64.90%
α and β simultaneously (Dr²Net)	68.13%

of action recognition [41]. The two coefficients α and β are updated every η iterations. A smaller value of η means that they are updated more frequently and in a smaller step. We can see from the table that the smaller η is, the higher performance we will obtain.

Updating policies are compared in Tab. 12 for the task of temporal action detection [44]. If we consider the values of α and β to be functions of training iterations, we can use different functions to represent different updating policies. We compare the linear updating policy in our Dr²Net to the exponential and logarithm policy in the table, and find that

Table 11. **Comparison of different values of the updating frequency η on action recognition with VideoMAE pretrained model [41].** The two coefficients α and β are updated every η iterations. Smaller η values means more frequent update, and it shows higher accuracy than larger η values.

η	2 iter	5 iter	20 iter	50 iter	100 iter
Top-5 Acc (%)	89.01	88.75	88.61	88.49	88.08

Table 12. **Comparison of different updating policies for α and β .** We consider the values of α and β as functions of training iterations, and experiment with the following three functions as different updating policies. We report mean average precision (mAP) at tIoU thresholds 0.5, 0.75, and 0.95, and average mAP. The linear function shows the highest accuracy.

Policy	0.5	0.75	0.95	Avg. mAP
Exponential	52.68%	35.93%	9.47%	35.46%
Logarithm	52.43%	36.49%	9.38%	35.69%
Linear (Dr²Net)	53.24%	36.97%	10.16%	36.27%

our linear policy leads to the best performance.

5. Conclusions

In this paper, we propose Dynamic Reversible Dual-Residual Networks (Dr²Net), a novel approach for finetuning pretrained models with significantly reduced memory usage. Dr²Net contains two types of residual connections, one maintaining the residual structure in the pretrained models, and the other introducing reversibility to enable clearing of intermediate activations from memory during training. We adopt a dynamic finetuning strategy that ensures a smooth transition from the non-reversible pretrained network to the reversible network. Evaluation across various tasks demonstrates that Dr²Net achieves performance comparable to conventional finetuning methods but with much lower memory requirements.

This work presents a practical solution for scenarios where downstream tasks are hindered by excessive memory consumption or restricted memory capacity. This includes applications involving large models, tasks dealing with high-resolution or high-dimensional data, and on-device learning environments. It could open avenues for future research in memory-efficient network architectures within the field of computer vision, as well as extending its implications to applications beyond computer vision, including natural language processing and audio analysis.

Acknowledgement. This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research through the Visual Computing Center (VCC) funding, as well as the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence (SDAIA-KAUST AI).

References

- [1] Juan Leon Alcazar, Moritz Cordes, Chen Zhao, and Bernard Ghanem. End-to-end active speaker detection. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1
- [2] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3D semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 4
- [4] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multimodal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 6, 7
- [5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. ActivityNet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1, 6
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 4
- [7] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016. 3
- [8] Feng Cheng and Gedas Bertasius. TALLFormer: Temporal action localization with long-memory transformer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 1, 2
- [10] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. 3
- [11] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 3
- [12] 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 16×16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021. 1, 2, 3, 6
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International conference on computer vision (ICCV)*, 2019. 2
- [14] Christoph Feichtenhofer, Yanghao Li, Kaiming He, et al. Masked autoencoders as spatiotemporal learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1
- [15] Kirill Gavrilyuk, Amir Ghodrati, Zhenyang Li, and Cees GM Snoek. Actor and action video segmentation from a sentence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5958–5966, 2018. 6
- [16] Aidan N Gomez, Mengye Ren, Raquel Urtasun, and Roger B Grosse. The reversible residual network: Backpropagation without storing activations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 1, 2, 3, 5, 7
- [17] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 6
- [18] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4D: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [19] 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 (CVPR)*, 2016. 3
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2022. 1, 2, 3, 6
- [21] Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pages 2722–2730. PMLR, 2019. 3
- [22] Jörn-Henrik Jacobsen, Arnold Smeulders, and Edouard Oyallon. i-RevNet: Deep invertible networks. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [23] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018. 3
- [24] Guohao Li, Matthias Müller, Bernard Ghanem, and Vladlen Koltun. Training graph neural networks with 1000 layers. In *International Conference on Machine Learning (ICML)*, 2021. 2
- [25] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. BMN: Boundary-matching network for temporal action proposal generation. In *Proceedings of The IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence

- Zitnick. Microsoft COCO: Common objects in context. In *Proceedings of the European conference on computer vision (ECCV)*, 2014. 6
- [27] Shuming Liu, Mengmeng Xu, Chen Zhao, Xu Zhao, and Bernard Ghanem. ETAD: A unified framework for efficient temporal action detection. *arXiv preprint arXiv:2205.07134*, 2022. 1, 2, 3
- [28] Shuming Liu, Chen-Lin Zhang, Chen Zhao, and Bernard Ghanem. End-to-end temporal action detection with 1b parameters across 1000 frames. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3
- [29] Yang Liu, Zhenyue Qin, Saeed Anwar, Pan Ji, Dongwoo Kim, Sabrina Caldwell, and Tom Gedeon. Invertible denoising network: A light solution for real noise removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2022. 1, 2, 3, 6
- [31] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 5, 6
- [32] Karttikeya Mangalam, Haoqi Fan, Yanghao Li, Chao-Yuan Wu, Bo Xiong, Christoph Feichtenhofer, and Jitendra Malik. Reversible vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10830–10840, 2022. 3
- [33] Karttikeya Mangalam, Haoqi Fan, Yanghao Li, Chao-Yuan Wu, Bo Xiong, Christoph Feichtenhofer, and Jitendra Malik. Reversible vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 5, 7
- [34] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. *arXiv preprint arXiv:1710.03740*, 2017. 2
- [35] Chong Mou, Youmin Xu, Jiechong Song, Chen Zhao, Bernard Ghanem, and Jian Zhang. Large-capacity and flexible video steganography via invertible neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [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, 2
- [37] Guocheng Qian, Abdullah Hamdi, Xingdi Zhang, and Bernard Ghanem. Image pretrained standard transformers for 3d point cloud understanding. In *International Conference on 3D Vision (3DV)*, 2024. 1, 2, 6, 7, 8
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning (ICML)*, 2021. 1, 2
- [39] Mery Ramazanova, Victor Escorcía, Fabian Caba Heilbron, Chen Zhao, and Bernard Ghanem. Owl (observe, watch, listen): Localizing actions in egocentric video via audiovisual temporal context. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2023. 3
- [40] Mattia Soldan, Alejandro Pardo, Juan León Alcázar, Fabian Caba Heilbron, Chen Zhao, Silvio Giancola, and Bernard Ghanem. Mad: A scalable dataset for language grounding in videos from movie audio descriptions. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [41] 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*, 35:10078–10093, 2022. 1, 2, 3, 6, 7, 8
- [42] Mengmeng Xu, Chen Zhao, David S Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3
- [43] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel Ni, and Heung-Yeung Shum. DINO: DETR with improved denoising anchor boxes for end-to-end object detection. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 2, 3, 6, 7, 8
- [44] Chen Zhao, Ali K Thabet, and Bernard Ghanem. Video self-stitching graph network for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 3, 6, 7, 8
- [45] Chen Zhao, Mery Ramazanova, Mengmeng Xu, and Bernard Ghanem. Segtad: Precise temporal action detection via semantic segmentation. *Proceedings of the European Conference on Computer Vision Workshop (ECCVW)*, 2022. 3
- [46] Chen Zhao, Shuming Liu, Karttikeya Mangalam, and Bernard Ghanem. Re²TAL: Rewiring pretrained video backbones for reversible temporal action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3, 4, 5, 6, 7
- [47] Andrew Zisserman, Joao Carreira, Karen Simonyan, Will Kay, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, et al. The Kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1, 2, 3