

# Omni-sourced Webly-supervised Learning for Video Recognition

Haodong Duan<sup>1</sup>, Yue Zhao<sup>1</sup>, Yuanjun Xiong<sup>2</sup>, Wentao Liu<sup>3</sup>, and Dahua Lin<sup>1</sup>

<sup>1</sup> The Chinese University of Hong Kong

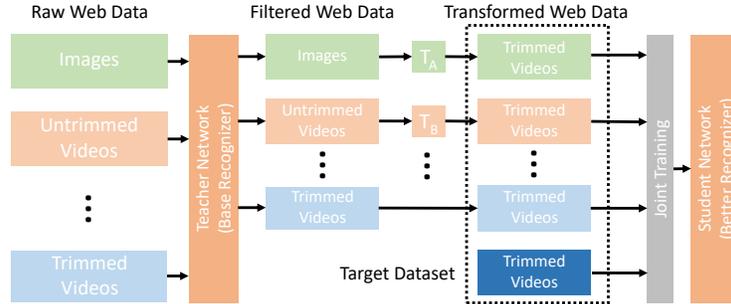
<sup>2</sup> Amazon AI

<sup>3</sup> SenseTime Research

**Abstract.** We introduce OmniSource, a novel framework for leveraging web data to train video recognition models. OmniSource overcomes the barriers between data formats, such as images, short videos, and long untrimmed videos for webly-supervised learning. First, data samples with multiple formats, curated by task-specific data collection and automatically filtered by a teacher model, are transformed into a unified form. Then a joint-training strategy is proposed to deal with the domain gaps between multiple data sources and formats in webly-supervised learning. Several good practices, including data balancing, resampling, and cross-dataset mixup are adopted in joint training. Experiments show that by utilizing data from multiple sources and formats, OmniSource is more data-efficient in training. With only 3.5M images and 800K minutes videos crawled from the internet without human labeling (less than 2% of prior works), our models learned with OmniSource improve Top-1 accuracy of 2D- and 3D-ConvNet baseline models by 3.0% and 3.9%, respectively, on the Kinetics-400 benchmark. With OmniSource, we establish new records with different pretraining strategies for video recognition. Our best models achieve **80.4%**, **80.5%**, and **83.6%** Top-1 accuracies on the Kinetics-400 benchmark respectively for training-from-scratch, ImageNet pre-training and IG-65M pre-training.

## 1 Introduction

Following the great success of representation learning in image recognition [22,39,15,17], recent years have witnessed great progress in video classification thanks to the development of stronger models [38,47,3,43] as well as the collection of larger-scale datasets [3,56,32,31]. However, labelling large-scale image datasets [37,59] is well known to be costly and time-consuming. It is even more difficult to do so for trimmed video recognition. The reason is that most online videos are *untrimmed*, *i.e.* containing numerous shots with multiple concepts, making it unavoidable to first go through the entire video and then manually cut it into informative video clips based on a specific query. Such procedure requires far more efforts than image annotation where a simple glance and click is needed. As a result, while the quantity of web videos grows exponentially over the past 3 years, the Kinetics dataset merely grows from 300K videos in 400 classes [19] to 650K in 700 classes [2], partially limiting the scaling-up of video architectures [3].



**Fig. 1: OmniSource Framework.** We first train a teacher network on the target dataset. Then, we use the teacher network to filter collected web data of different formats, to reduce noise and improve data quality. Specific transformations are conducted on the filtered out data corresponding to their formats. The target dataset and auxiliary web datasets are used for joint training of the student network

Instead of confining ourselves to the well-annotated trimmed videos, we move beyond by exploring the abundant visual data that are publicly available on the Internet in a more labor-saving way. These visual data are in various formats, including images, short video clips, and long videos. They capture the same visual world while exhibiting different advantages: *e.g.* images may be of higher quality and focus on distinctive moments; short videos may be edited by the user, therefore contain denser information; long videos may depict an event in multiple views. We transform the data in different formats into a unified form so that a single model can combine the best of both worlds.

Recent works [29,12] explore the possibility of pre-training from massive unlabeled web images or videos only with hashtags. However, they restrict the scope to the data of a single format. Also, these methods usually require billions of images to obtain a pre-trained 2D CNN model that is resilient to noise, which poses great costs and restricts its practicability. Besides, to take advantage of representation learned from large-scale images for videos, we have to take extra steps to transfer the 2D ConvNets to the 3D counterparts, either by inflating [3] or distillation [13], and then perform fine-tuning on the target dataset, which is tedious and may be suboptimal.

In this work, we propose a simple and unified framework for video classification while utilizing multiple sources of web data in different formats simultaneously. To enhance data efficiency, we propose task-specific data collection, *i.e.* obtaining topmost results using class labels as keywords on search engines, making the supervision most informative. Our framework consists of three steps: (1) We train one (or more) teacher network on the labeled dataset; (2) For each source of data collected, we apply the corresponding teacher network to obtain pseudo-labels and filter out irrelevant samples with low confidence; (3) We apply different transforms to convert each type of web data (*e.g.* images) to the target input format (*e.g.* video clips) and train the student network.

There are two main obstacles during joint training with the labeled dataset and unlabeled web datasets. First, possible domain gaps occur. For example, web images may focus more on objects and contain less motion blur than videos. Second, teacher filtering may lead to unbalanced data distribution across different classes. To mitigate the domain gap, we propose to balance the size of training batches between the labeled dataset and unlabeled web datasets and apply cross-dataset *mixup*. To cope with data imbalance, we try several resampling strategies. All these techniques contribute to the success of our approach.

Compared to the previous methods, our method excels at the following aspects: (1) It leverages a mixture of web data forms, including images, trimmed videos and untrimmed videos into one student network, aiming at an *omni-sourced* fashion. (2) It is *data-efficient*. Empirical results show that only 2M images, a significantly smaller amount compared to the total frame number of Kinetics (240K videos,  $\sim 70$ M frames), are needed to produce notable improvements (about 1%). For trimmed videos, the required amount is around 0.5M. In stark contrast, 65M videos are collected to obtain a noise-resilient pre-trained model in [12,49]. It is also noteworthy that our framework can also benefit from the massively weakly-supervised pre-training from billions of images or videos.

To sum up, our contributions are as follows:

- (1) We propose OmniSource, a simple and efficient framework for webly-supervised video classification, which can leverage web data in different formats.
- (2) We propose good practices for problems during joint training with omni-sourced data, include source-target balancing, resampling and cross-dataset mixup.
- (3) In experiments, our models trained by OmniSource achieve state-of-the-art performance on the Kinetics-400, for all pre-training strategies we tested.

## 2 Related work

**Webly-supervised learning** Leveraging information from the Internet, termed *webly-supervised learning*, has been extensively explored [25,50,14,11]. Divvala *et al* in [7] proposes to automatically learn models from online resources for visual concept discovery and image annotation. Chen *et al* reveals that images crawled from the Internet can yield superior results over the fully-supervised method [5]. For video classification, Ma *et al* proposes to use web images to boost action recognition models in [28] at the cost of manually filtering web action images. To free from additional human labor, efforts have been made to learn video concept detectors [51,26] or to select relevant frames from videos [10,41,52]. These methods are based on frames thus fail to consider the rich temporal dynamics of videos. Recent works [29,12] show that webly-supervised learning can produce better pre-training models with very large scale noisy data ( $\sim 10^9$  images and  $\sim 10^7$  videos). Being orthogonal to the pre-training stage, our framework works in a joint-training paradigm and is complementary to large-scale pre-training.

**Semi-supervised learning** Our framework works under the semi-supervised setting where labeled and unlabeled(web) data co-exist. Representative classical approaches include label propagation [61], self-training [36], co-training [1], and

graph networks [21]. Deep models make it possible to learn directly from unlabeled data via generative models [20], self-supervised learning [53], or consensus of multiple experts [54]. However, most existing methods are validated only on small scale datasets. One concurrent work [49] proposes to first train a student network with unlabeled data with pseudo-labels and then fine-tune it on the labeled dataset. Our framework, however, works on the two sources simultaneously, free from the pretrain-finetune paradigm and is more data-efficient.

**Distillation** According to the setting of knowledge distillation [16] and data distillation [35], given a set of manually labeled data, we can train a base model in the manner of supervised learning. The model is then applied to the unlabeled data or its transforms. Most of the previous efforts [35] are confined to the domain of images. In [13], Rohit *et al* proposes to distill spatial-temporal features from unlabeled videos with image-based teacher networks. Our framework is capable of distilling knowledge from multiple sources and formats within a single network.

**Domain Adaptation** Since web data from multiple sources are taken as input, domain gaps inevitably exist. Previous efforts [6,45,4] in domain adaptation focus on mitigating the data shift [34] in terms of data distributions. On the contrary, our framework focuses on adapting visual information in different formats (*e.g.* still images, long videos) into the same format (*i.e.* trimmed video clips).

**Video classification** Video analysis has long been tackled using hand-crafted feature [24,46]. Following the success of deep learning for images, video classification architectures have been dominated by two families of models, *i.e.* two-stream [38,47] and 3D ConvNets [3,44]. The former uses 2D networks to extract image-level feature and performs temporal aggregation [47,58,18] on top while the latter learns spatial-temporal features directly from video clips [44,8,43].

### 3 Method

#### 3.1 Overview

We propose a unified framework for omni-sourced webly-supervised video recognition, formulated in Sec. 3.2. The framework exploits web data of various forms (images, trimmed videos, untrimmed videos) from various sources (search engine, social media, video sharing platform) in an integrated way. Since web data can be very noisy, we use a teacher network to filter out samples with low confidence scores and obtain pseudo labels for the remaining ones (Sec. 3.4). We devise transformations for each form of data to make them applicable for the target task in Sec. 3.5. In addition, we explore several techniques to improve the robustness of joint training with web data in Sec. 3.6.

#### 3.2 Framework formulation

Given a target task (trimmed video recognition, *e.g.*) and its corresponding *target* dataset  $\mathcal{D}_{\mathcal{T}} = \{(\mathbf{x}_i, \mathbf{y}_i)\}$ , we aim to harness information from unlabeled web resources  $\mathcal{U} = \mathcal{U}_1 \cup \dots \cup \mathcal{U}_n$ , where  $\mathcal{U}_i$  refers to unlabeled data

in a specific source or format. **First**, we construct the pseudo-labeled dataset  $\widehat{\mathcal{D}}_i$  from  $\mathcal{U}_i$ . Samples with low confidence are dropped using a teacher model  $\mathcal{M}$  trained on  $\mathcal{D}_{\mathcal{T}}$ , and the remaining data are assigned with pseudo-labels  $\widehat{\mathbf{y}} = \text{PseudoLabel}(\mathcal{M}(\mathbf{x}))$ . **Second**, we devise appropriate transforms  $\mathcal{T}_i(\mathbf{x}) : \widehat{\mathcal{D}}_i \rightarrow \mathcal{D}_{\mathcal{A},i}$  to process data in a specific format (*e.g.* still images or long videos) into the data format (trimmed videos in our case) in the target task. We denote the union of  $\mathcal{D}_{\mathcal{A},i}$  to be the *auxiliary* dataset  $\mathcal{D}_{\mathcal{A}}$ . **Finally**, a model  $\mathcal{M}'$  (not necessarily the original  $\mathcal{M}$ ), can be jointly trained on  $\mathcal{D}_{\mathcal{T}}$  and  $\mathcal{D}_{\mathcal{A}}$ . In each iteration, we sample two mini-batches of data  $\mathcal{B}_{\mathcal{T}}, \mathcal{B}_{\mathcal{A}}$  from  $\mathcal{D}_{\mathcal{T}}, \mathcal{D}_{\mathcal{A}}$  respectively. The loss is a sum of cross entropy loss on both  $\mathcal{B}_{\mathcal{T}}$  and  $\mathcal{B}_{\mathcal{A}}$ , indicated by Eq 1.

$$\mathcal{L} = \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{B}_{\mathcal{T}}} \mathcal{L}(\mathcal{F}(\mathbf{x}; \mathcal{M}'), \mathbf{y}) + \sum_{\mathbf{x}, \widehat{\mathbf{y}} \in \mathcal{B}_{\mathcal{A}}} \mathcal{L}(\mathcal{F}(\mathbf{x}; \mathcal{M}'), \widehat{\mathbf{y}}) \quad (1)$$

For clarification, we compare our framework with some recent works on billion-scale webly-supervised learning in Table 1. OmniSource is capable of dealing with web data from multiple sources. It is designed to help a specific task, treats webly-supervision as co-training across multiple data sources instead of pre-training, thus is much more data-efficient. It is also noteworthy that our framework is orthogonal to webly-supervised pre-training [12].

**Table 1: Difference to previous works.** The notions follow Sec. 3.2:  $\mathcal{U}$  is the unlabeled web data,  $\mathcal{D}_{\mathcal{T}}$  is the target dataset.  $|\mathcal{U}|, |\mathcal{D}_{\mathcal{A}}|$  denotes the scale of web data and filtered auxiliary dataset

	Webly-supervised pretrain [29,12]	Web-scale semi-supervised [49]	OmniSource (Ours)
Procedure	<ol style="list-style-type: none"> <li>1. Train a model <math>\mathcal{M}</math> on <math>\mathcal{U}</math>.</li> <li>2. Fine-tune <math>\mathcal{M}</math> on <math>\mathcal{D}_{\mathcal{T}}</math>.</li> </ol>	<ol style="list-style-type: none"> <li>1. Train a model <math>\mathcal{M}</math> on <math>\mathcal{D}_{\mathcal{T}}</math>.</li> <li>2. Run <math>\mathcal{M}</math> on <math>\mathcal{U}</math> to pseudo-labeled <math>\widehat{\mathcal{D}}</math>.</li> <li>3. Train a student model <math>\mathcal{M}'</math> on <math>\widehat{\mathcal{D}}</math>.</li> <li>4. Fine-tune <math>\mathcal{M}'</math> on <math>\mathcal{D}_{\mathcal{T}}</math>.</li> </ol>	<ol style="list-style-type: none"> <li>1. Train one (or more) model <math>\mathcal{M}</math> on <math>\mathcal{D}_{\mathcal{T}}</math>.</li> <li>2. Run <math>\mathcal{M}</math> on <math>\bigcup_i \mathcal{U}_i</math> to pseudo-labeled <math>\bigcup_i \widehat{\mathcal{D}}_i</math>. (Samples under certain threshold are dropped.)</li> <li>3. Apply transforms <math>\mathcal{T}_i : \widehat{\mathcal{D}}_i \rightarrow \mathcal{D}_{\mathcal{A},i}</math>.</li> <li>4. Train model <math>\mathcal{M}'</math> (or <math>\mathcal{M}</math>) on <math>\mathcal{D}_{\mathcal{T}} \cup \mathcal{D}_{\mathcal{A}}</math>.</li> </ol>
$ \mathcal{U} $	3.5B images or 65M videos	1B images or 65M videos	$ \mathcal{U} $ : 13M images and 1.4M videos (0.4%~2%)
$ \mathcal{D}_{\mathcal{A}} $			$ \mathcal{D}_{\mathcal{A}} $ : 3.5M images and 0.8M videos (0.1%~1%)

### 3.3 Task-specific data collection

We use class names as keywords for data **crawling**, with no extra query expansion. For tag-based system like Instagram, we use automatic permutation and stemming<sup>4</sup> to generate tags. We crawl web data from various sources, including search engine, social media and video sharing platform. Because Google restricts the number of results for each query, we conduct multiple queries, each of which is restricted by a specific period of time. Comparing with previous works [29,12] which rely on large-scale web data with hashtags, our task-specific collection uses keywords highly correlated with labels, making the supervision stronger. Moreover, it reduces the required amount of web data by 2 orders of magnitude (*e.g.* from 65M to 0.5M videos on Instagram).

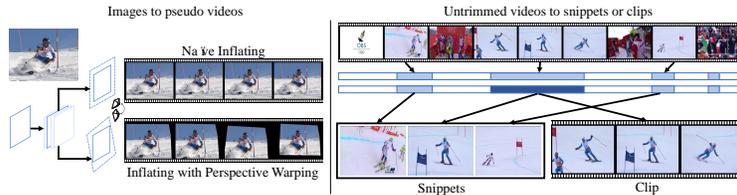
After data collection, we first remove invalid or corrupted data. Since web data may contain samples very similar to validation data, data **de-duplication**

<sup>4</sup> For example, “beekeeping” can be transformed to “beekeep”, and “keeping bee”.

is essential for a fair comparison. We perform content-based data de-duplication based on feature similarity. First, we extract frame-level features using an ImageNet-pretrained ResNet50. Then, we calculate the cosine similarity of features between the web data and target dataset and perform pairwise comparison after whitening. The average similarity among different crops of the same frame is used as the threshold. Similarity above it indicates suspicious duplicates. For Kinetics-400, we filter out 4,000 web images (out of 3.5M, 0.1%) and 400 web videos (out of 0.5M, 0.1%). We manually inspect a subset of them and find that less than 10% are real duplicates.

### 3.4 Teacher filtering

Data crawled from the web are inevitably noisy. Directly using collected web data for joint training leads to a significant performance drop (over 3%). To prevent irrelevant data from polluting the training set, we first train a teacher network  $\mathcal{M}$  on the target dataset and discard those web data with low confidence scores. For web images, we observe performance deterioration when deflating 3D teachers to 2D and therefore only use 2D teachers. For web videos, we find both applicable and 3D teachers outperform 2D counterparts consistently.



**Fig. 2: Transformations.** Left: Inflating images to clips, by replicating or inflating with perspective warping; Right: Extracting segments or clips from untrimmed videos, guided by confidence scores

### 3.5 Transforming to the target domain

**Web Images.** To prepare web images for video recognition training, we devise several ways to transform images into pseudo videos. The first naïve way is to replicate the image  $n$  times to form an  $n$ -frame clip. However, such clips may not be optimal since there is a visible gap between static clips and natural videos which visually change over time. Therefore, we propose to generate video clips from static images by viewing them with a moving camera. Given an image  $I$ , under the standard perspective projection model [9], an image with another perspective  $\tilde{I}$  can be generated by a homographic transform  $\mathcal{H}$  which is induced by a homographic matrix  $\mathbf{H} \in \mathbb{R}^{3 \times 3}$ , *i.e.*,  $\tilde{I} = \mathcal{H}(I) = \mathcal{F}(I; \mathbf{H})$ . To generate a clip  $J = \{J_1, \dots, J_N\}$  from  $I$ , starting from  $J_1 = I$ , we have

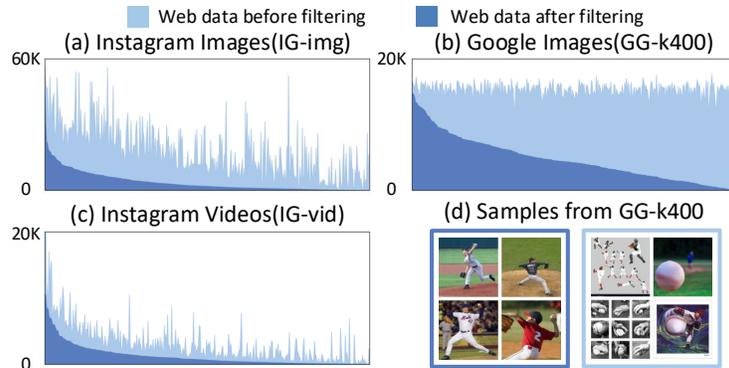
$$J_i = \mathcal{H}_i(J_{i-1}) = (\mathcal{H}_i \circ \mathcal{H}_{i-1} \circ \dots \circ \mathcal{H}_1)(I) \quad (2)$$

Each matrix  $\mathbf{H}_i$  is randomly sampled from a multivariate Gaussian distribution  $\mathcal{N}(\mu, \Sigma)$ , while the parameters  $\mu$  and  $\Sigma$  are estimated using maximum likelihood estimation on the original video source. Once we get pseudo videos, we can leverage web images for joint training with trimmed video datasets.

**Untrimmed Videos.** Untrimmed videos form an important part of web data. To exploit web untrimmed videos for video recognition, we adopt different transformations respectively for 2D and 3D architectures.

For 2D TSN, *snippets* sparsely sampled from the entire video are used as input. We first extract frames from the entire video at a low frame rate (1 FPS). A 2D teacher is used to get the confidence score of each frame, which also divides frames into positive ones and negative ones. In practice, we find that only using positive frames to construct snippets is a sub-optimal choice. Instead, combining negative frames and positive frames can form harder examples, results in better recognition performance. In our experiments, we use 1 positive frame and 2 negative frames to construct a 3-snippet input.

For 3D ConvNets, video *clips* (densely sampled continuous frames) are used as input. We first cut untrimmed videos into 10-second clips, then use a 3D teacher to obtain confidence scores. Only positive clips are used for joint training.



**Fig. 3: Web Data Distribution.** The inter-class distribution of three web datasets is visualized in (a,b,c), both before and after filtering. (d) gives out samples of filtered out images (cyan) and remained images (blue) for GG-K400. Teacher filtering successfully filters out lots of negative examples while making inter-class distribution more uneven

### 3.6 Joint training

Once web data are filtered and transformed into the same format of that in the target dataset  $\mathcal{D}_{\mathcal{T}}$ , we construct an *auxiliary* dataset  $\mathcal{D}_{\mathcal{A}}$ . A network can then be trained with both  $\mathcal{D}_{\mathcal{T}}$  and  $\mathcal{D}_{\mathcal{A}}$  using sum of cross-entropy loss in Eq. 1. As shown in Fig. 3, web data across classes are extremely unbalanced, especially after teacher filtering. Also there exists potential domain gap between  $\mathcal{D}_{\mathcal{T}}$  and  $\mathcal{D}_{\mathcal{A}}$ . To mitigate these issues, we enumerate several good practices as follows.

**Balance between target and auxiliary mini-batches.** Since the auxiliary dataset may be much larger than the target dataset and the domain gap may occur, the data ratio between target and auxiliary mini-batches is crucial for the final performance. Empirically,  $|\mathcal{B}_{\mathcal{T}}| : |\mathcal{B}_{\mathcal{A}}| = 2 : 1 \sim 1 : 1$  works reasonably well.

**Resampling strategy.** Web data are extremely unbalanced, especially after teacher filtering (see Fig 3). To alleviate this, we explore several sampling policies: (1) sampling from a clipped distribution: classes whose samples exceeds threshold  $N_c$  are clipped; (2) sampling from distribution modified by a power law: the probability of choosing class with  $N$  samples is proportional to  $N^p$  ( $p \in (0, 1)$ ). We find that (2) parameterized by  $p = 0.2$  is generally a better practice.

**Cross-dataset mixup.** Mixup [55] is a widely used strategy in image recognition. It uses convex combinations of pairs of examples and their labels for training, thus improving the generalization of deep neural networks. We find that technique also works for video recognition. When training teacher networks on  $\mathcal{D}_{\mathcal{T}}$  only, we use the linear combination of two clip-label pairs as training data, termed as *intra-dataset* mixup. When both target and auxiliary datasets are used, the two pairs are samples randomly chosen from both datasets, termed as *cross-dataset* mixup. Mixup works fairly well when networks are trained from scratch. For fine-tuning, the performance gain is less noticeable.

## 4 Datasets

In this section, we introduce the datasets on which experiments will be conducted. Then we go through different sources from which web data are collected.

### 4.1 Target datasets

**Kinetics-400** The Kinetics dataset [3] is one of the largest video datasets. We use the version released in 2017 which contains 400 classes and each category has more than 400 videos. In total, it has around 240K, 19K, and 38K videos for training, validation and testing subset respectively. In each video, a 10-second clip is annotated and assigned a label. These 10-second clips constitute the data source for the default supervised learning setting, which we refer to **K400-tr**. The rest part of training videos is used to mimic untrimmed videos sourced from the Internet which we refer to **K400-untr**.

**Youtube-car** Youtube-car [60] is a fine-grained video dataset with 10K training and 5K testing videos of 196 types of cars. The videos are untrimmed, last several minutes. Following [60], the frames are extracted from videos at 4 FPS.

**UCF101** UCF101 [40] is a small scale video recognition dataset, which has 101 classes and each class has around 100 videos. We use the official split-1 in our experiments, which has about 10K and 3.6K videos for training and testing.

### 4.2 Web sources

We collect web images and videos from various sources including search engines, social medias and video sharing platforms.

**GoogleImage** GoogleImage is a search engine based web data source for Kinetics-400, Youtube-car and UCF101. We query each class name in the target dataset on Google to get related web images. We crawl 6M, 70K, 200K URLs for Kinetics-400, Youtube-car and UCF101 respectively. After data cleaning and teacher filtering, about 2M, 50K, 100K images are used for training on these three datasets. We denote the three datasets as **GG-k400**, **GG-car**, and **GG-UCF** respectively.

**Instagram** Instagram is a social media based web data source for Kinetics-400. It consists of InstagramImage and InstagramVideo. We generate several tags for each class in Kinetics-400, resulting in 1,479 tags and 8.7M URLs. After removing corrupted data and teacher filtering, about 1.5M images and 500K videos are used for joint training, denoted as **IG-img** and **IG-vid**. As shown in Fig 3, IG-img is significantly unbalanced after teacher filtering. Therefore, in the coming experiments, IG-img is used in combination with GG-k400.

**YoutubeVideo** YoutubeVideo is a video sharing platform based web data source for Youtube-car. We crawl 28K videos from youtube by querying class names. After de-duplicating (remove videos in the original Youtube-car dataset) and teacher filtering, 17K videos remain, which we denote as **YT-car-17k**.

## 5 Experiments

### 5.1 Video architectures

We mainly study two families of video classification architectures, namely Temporal Segment Networks [47] and 3D ConvNets [3], to verify the effectiveness of our design. Unless specified, we use ImageNet-pretrained models for initialization. We conduct all experiments using MMAAction [57].

**2D TSN** Different from the original setting in [47], we choose ResNet-50 [15] to be the backbone, unless otherwise specified. The number of segments is set to be 3 for Kinetics/UCF-101 and 4 for Youtube-car, respectively.

**3D ConvNets** For 3D ConvNet, we use the SlowOnly architecture proposed in [8] in most of our experiments. It takes 64 consecutive frames as a video clip and sparsely samples 4/8 frames to form the network input. Different initialization strategies are explored, including training from scratch and fine-tuning from a pre-trained model. Besides, more advanced architecture like Channel Separable Network [43] and more powerful pre-training (IG-65M [12]) is also explored.

### 5.2 Verifying the efficacy of OmniSource

We verify our framework’s efficacy by examining several questions.

**Why do we need teacher filtering and are search results good enough?**

Some may question the necessity of a teacher network for filtering under the impression that a modern search engine might have internally utilized a visual recognition model, possibly trained on massively annotated data, to help generate the search results. However, we argue that web data are inherently noisy and we observe nearly half of the returned results are irrelevant. More quantitatively, 70% - 80% of the web data are rejected by the teacher. On the other

hand, we conduct an experiment without teacher filtering. Directly using collected web data for joint training leads to a significant (over 3%) performance drop on TSN. This reveals that teacher filtering is necessary to help retain the useful information from the crawled web data while eliminating the useless.

**Does every data source contribute?** We explore the contribution of different source types: images, trimmed videos and untrimmed videos. For each data source, we construct auxiliary dataset and use it for joint training with  $K400-tr$ . Results in Table 2 reveal that every source contributes to improving accuracy on the target task. When combined, the performance is further improved.

**Table 2: Every source contributes.** We find that each source contributes to the target task. With all sources combined (we intuitively set the ratio as:  $K400-tr : Web-img : IG-vid : K400-untr = 2 : 1 : 1 : 1$ ), the improvement can be more considerable. The conclusion holds for both 2D TSN and 3D ConvNets (Format: Top-1 Acc/ Top-5 Acc)

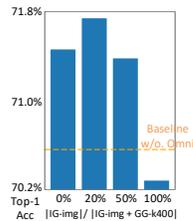
Arch/Dataset	K400-tr	+GG-k400	+GG&IG-img	+IG-vid	+K400-untr	+ All
TSN-3seg R50	70.6/89.4	71.5/89.5	72.0/90.0	72.0/90.3	71.7/89.6	73.6/91.0
SlowOnly 4x16,R50	73.8/90.9	74.5/91.4	75.2/91.6	75.2/91.7	74.5/91.1	76.6/92.5

For images, when the combination of  $GG-k400$  and  $IG-img$  is used, the Top-1 accuracy increases around 1.4%. For trimmed videos, we focus on  $IG-vid$ . Although being extremely unbalanced,  $IG-vid$  still improves Top-1 accuracy by over 1.0% in all settings. For untrimmed videos, we use the untrimmed version of Kinetics-400 ( $K400-untr$ ) as the video source and find it also works well.

**Do multiple sources outperform a single source?** Seeing that web data from multiple sources can jointly contribute to the target dataset, we wonder if multiple sources are still better than a single source with the same budget. To verify this, we consider the case of training TSN on both  $K400-tr$  and  $\mathcal{D}_A = GG-k400 + IG-img$ . We fix the scale of auxiliary dataset to be that of  $GG-k400$  and vary the ratio between  $GG-k400$  and  $IG-img$  by replacing images from  $GG-k400$  with those in  $IG-img$ . From Fig. 4, we observe an improvement of 0.3% without increasing  $|\mathcal{D}_A|$ , indicating that multiple sources provide complementary information by introducing diversity.

**Does OmniSource work with different architectures?** We further conduct experiments on a wide range of architectures and obtain the results in Table 3. For TSN, we use EfficientNet-B4 [42] instead as the backbone, on which OmniSource improves Top-1 accuracy by 1.9%. For 3D-ConvNets, we conduct experiments on the SlowOnly-8x8-ResNet101 baseline, which takes longer input and has a larger backbone. Our framework also works well in this case, improving the Top-1 accuracy from 76.3% to **80.4%** when training from scratch, from 76.8% to **80.5%** with ImageNet pretraining. The improvement on larger networks is higher, suggesting that deeper networks are more prone to suffering from the scarcity of video data and OmniSource can alleviate this.

Arch	Backbone	Pretrain	w/o. Omni	w/. Omni	$\Delta$
TSN-3seg	ResNet50	ImageNet	70.6 / 89.4	73.6 / 91.0	+3.0 / +1.6
TSN-3seg	ResNet50	IG-1B	73.1 / 90.4	75.7 / 91.9	+2.6 / +1.5
TSN-3seg	Efficient-b4	ImageNet	73.3 / 91.0	75.2 / 92.0	+1.9 / +1.0
SlowOnly-4x16	ResNet50	-	72.9 / 90.9	76.8 / 92.5	+3.9 / +1.6
SlowOnly-4x16	ResNet50	ImageNet	73.8 / 90.9	76.6 / 92.5	+2.8 / +1.6
SlowOnly-8x8	ResNet101	-	76.3 / 92.6	80.4 / 94.4	+4.1 / +1.8
SlowOnly-8x8	ResNet101	ImageNet	76.8 / 92.8	80.5 / 94.4	+3.7 / +1.6
irCSN-32x2	irCSN-152	IG-65M	82.6 / 95.3	83.6 / 96.0	+1.0 / +0.7



**Table 3: Improvement under various experiment configurations.** OmniSource is extensively tested on various architectures with various pretraining strategies. The improvement is significant in **ALL** tested choices. Even for the SOTA setting, which uses 65M web videos for pretraining, OmniSource still improves the Top-1 accuracy by 1.0% (Format: Top-1 / Top-5 Acc)

**Is OmniSource compatible with different pre-training strategies?** As discussed, OmniSource alleviates the data-hungry issue by utilizing auxiliary data. One natural question is: how does it perform when training 3D networks from scratch? Can we simply drop ImageNet pretraining in pursuit of a more straightforward training policy? Indeed, we find that OmniSource works fairly well under this setting and interestingly the performance gain is more significant than fine-tuning. For example, SlowOnly-(4x16, R50) increases the Top-1 accuracy by 3.9% when training from scratch while fine-tuning only increases by 2.8%. The model trained from scratch beats the fine-tuned counterpart by 0.2% with OmniSource though being 0.9% lower with only K400-tr. Similar results can be observed for SlowOnly-(8x8, R101). With large-scale webly supervised pretraining, OmniSource still leads to significant performance improvement (+2.6% Top-1 for TSN-3seg-R50, +1.0% Top-1 for irCSN-32x2).

**Fig. 4: Multi-source is better.** Mixed sources leads to better performance with a constrained number of web images

Arch	UCF101-split1		HMDB51-split1	
	w/o. Omni	w/. Omni	w/o. Omni	w. Omni
TSN-3seg R50[FT]	91.5	93.3	63.5	65.9
SlowOnly 4x16, R50[FT]	94.7	96.0	69.4	70.7
SlowOnly 4x16, R50[SC]	94.1	96.0	65.8	71.0



**Table 4: OmniSource features transfer well.** We finetune on UCF101 and HMDB51 with the K400-tr pretrained weight. Pretraining with OmniSource improves the performance significantly

**Fig. 5: Confusing pairs improved by OmniSource.** The original accuracy and change are denoted in black and in color

**Do features learned by OmniSource transfer to other tasks?** Although OmniSource is designed for a target video recognition task, the learned features also transfer well to other video recognition tasks. To evaluate the transfer capability, we finetune the learned model on two relatively smaller datasets: UCF101 [40] and HMDB51 [23]. Table 4 indicates that on both benchmarks, pretraining with OmniSource leads to significant performance improvements.

Following standard evaluation protocol, SlowOnly-8x8-R101 achieves 97.3% Top-1 accuracy on UCF101, 79.0% Top-1 accuracy on HMDB51 with RGB input. When combined with optical flow, it achieves **98.6%** and **83.8%** Top-1 accuracy on UCF101 and HMDB51, which is the new state-of-the-art. More results on transfer learning are provided in the supplementary material.

**Does OmniSource work in different target domains?** Our framework is also effective and efficient in various domains. For a fine-grained recognition benchmark called **Youtube-car**, we collect 50K web images (GG-car) and 17K web videos (YT-car-17k) for training. Table 5 shows that the performance gain is significant: 5% in both Top-1 accuracy and mAP. On **UCF-101**, we train a two-stream TSN network with BNInception as the backbone. The RGB stream is trained either with or without GG-UCF. The results are listed in Table 6. The Top-1 accuracy of the RGB stream improves by 2.7%. When fused with the flow stream, there is still an improvement of 1.1%.

**Where does the performance gain come from?** To find out why web data help, we delve deeper into the collected web dataset and analyze the improvement on individual classes. We choose TSN-3seg-R50 trained either with or without GG-k400, where the improvement is 0.9% on average. We mainly focus on the confusion pairs that web images can improve. We define the *confusion score* of a class pair as  $s_{ij} = (n_{ij} + n_{ji}) / (n_{ij} + n_{ji} + n_{ii} + n_{jj})$ , where  $n_{ij}$  denotes the number of images whose ground-truth are class  $i$  while being recognized as class  $j$ . Lower confusion score denotes better discriminating power between the two classes. We visualize some confusing pairs in Fig 5. We find the improvement can be mainly attributed to two reasons: (1) Web data usually focus on key objects of action. For example, we find that in those pairs with the largest confusion score reduction, there exist pairs like “drinking beer” vs. “drinking shots”, and “eating hotdog” vs. “eating chips”. Training with web data leads to better object recognition ability in some confusing cases. (2) Web data usually include discriminative poses, especially for those actions which last for a short time. For example, “rock scissors paper” vs. “shaking hands” has the second-largest confusion score reduction. Other examples including “sniffing”- “headbutting”, “break dancing”- “robot dancing”, etc.

**Table 7:** Comparisons with Kinetics-400 state-of-the-art

**Table 5:** Youtube-car

Setting	Top-1	mAP
Baseline	77.05	71.95
+GG-car	80.96	77.05
+YT-car-17k	81.68	78.61
+ [GG-] + [YT-]	81.95	78.67

**Table 6:** UCF-101

Setting	+ Flow	Top-1
Baseline		86.04
+ GG-UCF		88.74
Baseline	✓	93.47
+ GG-UCF	✓	94.58

Method	backbone	pretrain	Top-1	Top-5
TSN-7seg [47]	Inception-v3	ImageNet	73.9	91.1
TSM-8seg [27]	ResNet50	ImageNet	72.8	N/A
TSN-3seg (Ours)	ResNet50	ImageNet	73.6	91.0
TSN-3seg (Ours)	Efficient-b4	ImageNet	<b>75.2</b>	<b>92.0</b>
SlowOnly-8x8 [8]	ResNet101	-	75.9	N/A
SlowFast-8x8 [8]	ResNet101	-	77.9	93.2
SlowOnly-8x8 (Ours)	ResNet101	-	<b>80.4</b>	<b>94.4</b>
I3D-64x1 [3]	Inception-V1	ImageNet	72.1	90.3
NL-128x1 [48]	ResNet101	ImageNet	77.7	93.3
SlowFast-8x8 [8]	ResNet101	ImageNet	77.9	93.2
LGD-3D (RGB) [33]	ResNet101	ImageNet	79.4	94.4
STDFB [30]	ResNet152	ImageNet	78.8	93.6
SlowOnly-8x8 (Ours)	ResNet101	ImageNet	<b>80.5</b>	<b>94.4</b>
irCSN-32x2 [12]	irCSN-152	IG-65M	82.6	95.3
irCSN-32x2 (Ours)	irCSN-152	IG-65M	<b>83.6</b>	<b>96.0</b>

### 5.3 Comparisons with state-of-the-art

In Table 7, we compare OmniSource with current state-of-the-art on Kinetics-400. For 2D ConvNets, we obtain competitive performance with fewer segments and lighter backbones. For 3D ConvNets, considerable improvement is achieved for all pre-training settings with OmniSource applied. With IG-65M pre-trained irCSN-152, OmniSource achieves **83.6%** Top-1 accuracy, an absolute improvement of 1.0% with only 1.2% relatively more data, establishing a new record.

**Table 8:** Different ways to transform images into video clips. Still inflation is a strong baseline, while agnostic perspective warping performs best

Inflation	Top-1	Top-5
N/A	73.8	90.9
replication (still)	74.1	91.2
translation (random)	73.7	90.9
translation (constant)	73.8	90.8
perspective warp [spec]	74.4	91.3
perspective warp [agno]	74.5	91.4

**Table 9:** Mixup technique can be beneficial to the model performance, both for intra- and cross-dataset cases. However, it works only when the model is trained from scratch

Pretraining	w. mixup	w.GG-img	Top-1	Top-5
ImageNet			73.8	90.9
ImageNet	✓		73.6	91.1
None			72.9	90.9
None	✓		73.3	90.9
None		✓	74.1	91.0
None	✓	✓	74.4	91.4

### 5.4 Validating the good practices in OmniSource

We conduct several ablation experiments on techniques we introduced. The target dataset is K400-tr and the auxiliary dataset is GG-k400 unless specified.

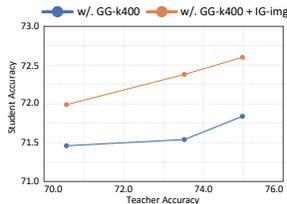
**Transforming images to video clips.** We compare different ways to transform web images into clips in Table 8. Naïvely replicating still image brings limited improvement (0.3%). We then apply translation with randomized or constant speed to form pseudo clips. However, the performance deteriorates slightly, suggesting that translation cannot mimic the camera motion well. Finally, we resort to perspective warping to hallucinate camera motion. Estimating class-agnostic distribution parameters is slightly better, suggesting that all videos might share similar camera motion statistics.

**Cross-Dataset mixup** In Table 9, we find that mixup is effective for video recognition in both intra- and cross-dataset cases when the model is trained from scratch. The effect is unclear for fine-tuning. In particular, mixup can lead to 0.4% and 0.3% Top-1 accuracy improvement for intra- and inter-dataset cases.

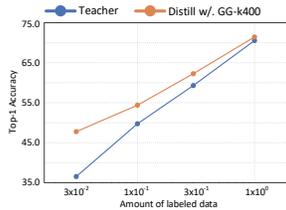
**Impact of teacher choice.** Since both teacher and student networks can be 2D or 3D ConvNets, there are 4 possible combinations for teacher network choosing. For images, deflating 3D ConvNets to 2D yields a dramatic performance drop. Therefore, we do not use 3D ConvNet teachers for web images. For videos, however, 3D ConvNets lead to better filtering results comparing to its 2D counterpart. To examine the effect of different teachers, we fix the student model to be a ResNet-50 and vary the choices of teacher models (ResNet-50, EfficientNet-b4, and the ensemble of ResNet-152 and EfficientNet-b4). Consistent improvement is observed against the baseline (70.6%). The student accuracy increases when a better teacher network is used. It also holds for 3D ConvNets on web videos.

**Effectiveness when labels are limited.** To validate the effectiveness with limited labeled data, we construct 3 subsets of K400-tr with a proportion of

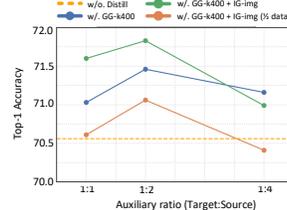
3%, 10%, and 30% respectively. We rerun the entire framework including data filtering with a weaker teacher. The final results on the validation set of `K400-tr` is shown in Fig 7. Our framework consistently improves the performance as the percentage of labeled videos varies. Particularly, the gain is more significant when data are scarce, *e.g.* a relative increase of over 30% with 3% labeled data.



**Fig. 6:** Better teachers leads to better students



**Fig. 7:** OmniSource on subsets of Kinetics



**Fig. 8:** Accuracy with different ratios  $|\mathcal{B}_T| : |\mathcal{B}_A|$

### Balancing between the target and auxiliary dataset.

We tune the ratio between the batch size of the target dataset  $|\mathcal{B}_T|$  and the auxiliary dataset  $|\mathcal{B}_A|$  and obtain the accuracy on Fig 8. We test 3 scenarios: (1) the original GG-k400, clarified in Sec 4.2; (2) [GG+IG]-k400, the union of GG-k400 and IG-img; (3) [GG+IG]-k400-half which is the half of (2). We observe that the performance gain is robust to the choice of  $|\mathcal{B}_T|/|\mathcal{B}_A|$  in most cases. However, with less auxiliary data, the ratio has to be treated more carefully. For example, smaller  $|\mathcal{D}_A|$  but larger  $|\mathcal{B}_A|$  may cause overfitting auxiliary samples and hurt the overall result.

**Resampling strategies.** The target dataset is usually balanced across classes. The nice property doesn't necessarily hold for the auxiliary dataset. Thus we propose several resampling strategies. From Table 10, we see that simple techniques to tailor the distribution into a more balanced one yield nontrivial improvements.

**Table 10: Resampling strategies.** Simple resampling strategies lead to nontrivial improvement

Strategy	Top-1/5
None (original)	71.5/89.5
Clipped ( $N_c = 5000$ )	71.9/90.0
Power ( $\sim N^p, p = 0.5$ )	71.8/89.7
Power ( $\sim N^p, p = 0.2$ )	72.0/90.0

## 6 Conclusion

In this work, we propose OmniSource, a simple yet effective framework for webly-supervised video recognition. Our method can utilize web data from multiple sources and formats by transforming them into a same format. In addition, our task-specific data collection is more data-efficient. The framework is applicable to various video tasks. Under all settings of pretraining strategies, we obtain state-of-the-art performance on multiple benchmarks.

**Acknowledgment** This work is partially supported by the SenseTime Collaborative Grant on Large-scale Multi-modality Analysis (CUHK Agreement No. TS1610626 & No. TS1712093), the General Research Fund (GRF) of Hong Kong (No. 14203518 & No. 14205719), and Innovation and Technology Support Program (ITSP) Tier 2, ITS/431/18F.

## References

1. Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, pages 92–100. ACM, 1998. [3](#)
2. Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. [1](#)
3. Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. [1](#), [2](#), [4](#), [8](#), [9](#), [12](#)
4. Min-Hung Chen, Zsolt Kira, Ghassan AlRegib, Jaekwon Yoo, Ruxin Chen, and Jian Zheng. Temporal attentive alignment for large-scale video domain adaptation. In *ICCV*, pages 6321–6330, 2019. [4](#)
5. Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*, pages 1431–1439, 2015. [3](#)
6. Gabriela Csurka. A comprehensive survey on domain adaptation for visual applications. In *Domain adaptation in computer vision applications*, pages 1–35. Springer, 2017. [4](#)
7. Santosh K Divvala, Ali Farhadi, and Carlos Guestrin. Learning everything about anything: Webly-supervised visual concept learning. In *CVPR*, pages 3270–3277, 2014. [3](#)
8. Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. [4](#), [9](#), [12](#)
9. David A Forsyth and Jean Ponce. *Computer vision: a modern approach*. Prentice Hall Professional Technical Reference, 2002. [6](#)
10. Chuang Gan, Chen Sun, Lixin Duan, and Boqing Gong. Webly-supervised video recognition by mutually voting for relevant web images and web video frames. In *ECCV*, pages 849–866. Springer, 2016. [3](#)
11. Chuang Gan, Ting Yao, Kuiyuan Yang, Yi Yang, and Tao Mei. You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images. In *CVPR*, pages 923–932, 2016. [3](#)
12. Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *CVPR*, pages 12046–12055, 2019. [2](#), [3](#), [5](#), [9](#), [12](#)
13. Rohit Girdhar, Du Tran, Lorenzo Torresani, and Deva Ramanan. Distinit: Learning video representations without a single labeled video. In *ICCV*, pages 852–861, 2019. [2](#), [4](#)
14. Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. [3](#)
15. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [1](#), [9](#)
16. Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. [4](#)
17. Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, pages 4700–4708, 2017. [1](#)
18. Noureldien Hussein, Efstratios Gavves, and Arnold WM Smeulders. Timeception for complex action recognition. In *CVPR*, pages 254–263, 2019. [4](#)
19. Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. [1](#)
20. Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *NeurIPS*, pages 3581–3589, 2014. [4](#)

21. Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 4
22. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1097–1105, 2012. 1
23. Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE, 2011. 11
24. Ivan Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005. 4
25. Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pages 5447–5456, 2018. 3
26. Junwei Liang, Lu Jiang, Deyu Meng, and Alexander G Hauptmann. Learning to detect concepts from webly-labeled video data. In *IJCAI*, pages 1746–1752, 2016. 3
27. Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. 12
28. Shugao Ma, Sarah Adel Bargal, Jianming Zhang, Leonid Sigal, and Stan Sclaroff. Do less and achieve more: Training cnns for action recognition utilizing action images from the web. *Pattern Recognition*, 68:334–345, 2017. 3
29. Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, pages 181–196, 2018. 2, 3, 5
30. Brais Martinez, Davide Modolo, Yuanjun Xiong, and Joseph Tighe. Action recognition with spatial-temporal discriminative filter banks. In *ICCV*, pages 5482–5491, 2019. 12
31. Antoinette Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640, 2019. 1
32. Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Yan Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 1
33. Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *CVPR*, pages 12056–12065, 2019. 12
34. Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. The MIT Press, 2009. 4
35. Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *CVPR*, pages 4119–4128, 2018. 4
36. Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. *WACV/MOTION*, 2, 2005. 3
37. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 1
38. Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, pages 568–576, 2014. 1, 4
39. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1

40. Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 8, 11
41. Chen Sun, Sanketh Shetty, Rahul Sukthankar, and Ram Nevatia. Temporal localization of fine-grained actions in videos by domain transfer from web images. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 371–380. ACM, 2015. 3
42. Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114, 2019. 10
43. Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, pages 5552–5561, 2019. 1, 4, 9
44. Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018. 4
45. Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 4
46. Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, pages 3551–3558, 2013. 4
47. Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1, 4, 9, 12
48. Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018. 12
49. I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546*, 2019. 3, 4, 5
50. Jufeng Yang, Xiaoxiao Sun, Yu-Kun Lai, Liang Zheng, and Ming-Ming Cheng. Recognition from web data: A progressive filtering approach. *IEEE Transactions on Image Processing*, 27(11):5303–5315, 2018. 3
51. Guangnan Ye, Yitong Li, Hongliang Xu, Dong Liu, and Shih-Fu Chang. Eventnet: A large scale structured concept library for complex event detection in video. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 471–480. ACM, 2015. 3
52. Serena Yeung, Vignesh Ramanathan, Olga Russakovsky, Liyue Shen, Greg Mori, and Li Fei-Fei. Learning to learn from noisy web videos. In *CVPR*, pages 5154–5162, 2017. 3
53. Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *ICCV*, pages 1476–1485, 2019. 4
54. Xiaohang Zhan, Ziwei Liu, Junjie Yan, Dahua Lin, and Chen Change Loy. Consensus-driven propagation in massive unlabeled data for face recognition. In *ECCV*, pages 568–583, 2018. 4
55. Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 8
56. Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *ICCV*, pages 8668–8678, 2019. 1
57. Yue Zhao, Haodong Duan, Yuanjun Xiong, and Dahua Lin. MMAAction. <https://github.com/open-mmlab/mmaaction>, 2019. 9
58. Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, pages 803–818, 2018. 4

59. Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017. [1](#)
60. Chen Zhu, Xiao Tan, Feng Zhou, Xiao Liu, Kaiyu Yue, Errui Ding, and Yi Ma. Fine-grained video categorization with redundancy reduction attention. In *ECCV*, pages 136–152, 2018. [8](#)
61. Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. *CMU CALD tech report CMU-CALD-02-107*, 2002. [3](#)