

Cross-Model Pseudo-Labeling for Semi-Supervised Action Recognition

Yinghao Xu¹ Fangyun Wei³ Xiao Sun³ Ceyuan Yang¹
Yujun Shen¹ Bo Dai² Bolei Zhou¹ Stephen Lin³

¹The Chinese University of Hong Kong ²S-Lab, Nanyang Technological University ³Microsoft Research Asia
{xy119, yc019, sy116, bzhou}@ie.cuhk.edu.hk bo.dai@ntu.edu.sg {fawe, xias, stevelin}@microsoft.com

Abstract

*Semi-supervised action recognition is a challenging but important task due to the high cost of data annotation. A common approach to this problem is to assign unlabeled data with pseudo-labels, which are then used as additional supervision in training. Typically in recent work, the pseudo-labels are obtained by training a model on the labeled data, and then using confident predictions from the model to teach itself. In this work, we propose a more effective pseudo-labeling scheme, called Cross-Model Pseudo-Labeling (CMPL). Concretely, we introduce a lightweight auxiliary network in addition to the primary backbone, and ask them to predict pseudo-labels for each other. We observe that, due to their different structural biases, these two models tend to learn complementary representations from the same video clips. Each model can thus benefit from its counterpart by utilizing cross-model predictions as supervision. Experiments on different data partition protocols demonstrate the significant improvement of our framework over existing alternatives. For example, CMPL achieves 17.6% and 25.1% Top-1 accuracy on Kinetics-400 and UCF-101 using only the RGB modality and 1% labeled data, outperforming our baseline model, FixMatch [17], by 9.0% and 10.3%, respectively.*¹

1. Introduction

The rapid development of deep learning has led to great success in action recognition. In the standard supervised learning protocol, a considerable number of annotated videos is needed but difficult to acquire in practice. On the other hand, about 500 hours of video is uploaded to YouTube every minute worldwide, providing a tremendous amount of unlabeled data. Leveraging such unlabeled videos for semi-supervised learning could thus be of great benefit for action recognition.

¹Project page is at <https://justimyhxu.github.io/projects/cmpl/>.

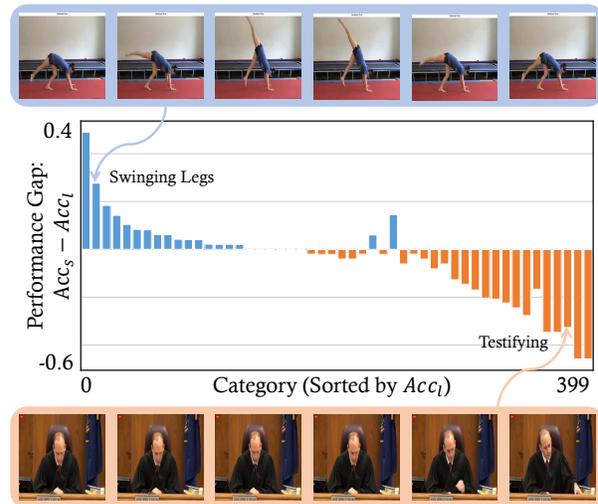


Figure 1. **Category-wise performance gap between small and large networks** under the supervised training setting given 1% labeled videos in Kinetics-400 [2]. Acc_s and Acc_l denote the accuracy of the small (3D-ResNet50 \times 1/4) and the large (3D-ResNet50) network, respectively. For categories on which the large network performs poorly (*i.e.*, the left half of the figure), the small network behaves better even with much lower model capacity. Concretely, the small network tends to perform well on classes with stronger temporal dynamics, *i.e.*, “Swinging Legs”, while the large network better recognizes actions mainly characterized by spatial information, *i.e.*, “Testifying”. See Sec. 4.4 for further discussion.

To gain supervision from unlabeled data, a common practice is to assign pseudo-labels to these data and treat them as “ground-truth” for training [11, 17, 19, 26]. Specifically, existing approaches train a model on labeled data and then use it to predict the unlabeled videos. If the confidence in a prediction is high enough, the prediction will be taken as the pseudo-label of the video, to be used in network training henceforth. The quantity and quality of pseudo-labels therefore has a significant impact in the current learning scheme. However, the limited discriminative power derived from a small amount of labeled data leads to inadequate pseudo-labels and limits the gain from unlabeled data.

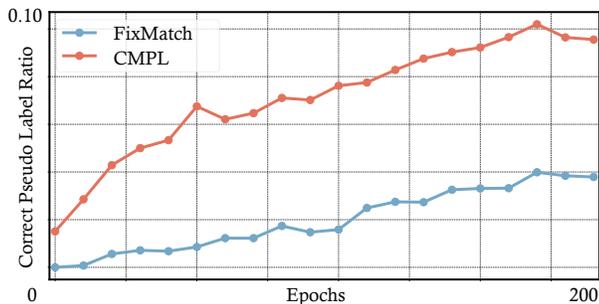


Figure 2. **Ratio of correct pseudo-labels.** To compare the quality and quantity of pseudo labels generated by CMPL and FixMatch, we plot the ratio of correct pseudo-labels relative to the total number of videos in the unlabeled dataset. Both models are trained with 1% labeled Kinetics-400 [2] (the others are unlabeled). Our CMPL is shown to more effectively provide high-quality pseudo-labels than FixMatch.

To better capitalize on unlabeled videos, we present a pseudo-labeling approach based on complementary representations at the model level. We observe that models of different scales exhibit markedly different behaviors in regards to category-wise performance, due to their different structural biases. As shown in Fig. 1 and studied in Sec. 4.4, a small model, despite its lower capacity, can achieve notable improvements over a large model on certain categories. In particular, it better captures temporal dynamics in recognizing actions, while a large model tends to better learn spatial semantics for distinguishing different action instances. This indicates that the two models differ in what they learn and therefore can complement each other in pseudo-labeling, where one model can more successfully generate pseudo-labels for some categories while the second model can be more effective on others.

Based on this observation, we propose *Cross-Model Pseudo-Labeling (CMPL)*, where the primary backbone is supplemented by a lightweight auxiliary network with a different structure and fewer channels than the backbone. This difference in architecture leads to a different representation of the input data that complements that of the primary backbone. Then, given an unlabeled video clip, we borrow a confident prediction from the auxiliary network as the pseudo-label for the primary backbone, and vice versa. As these two models have their own strengths, a greater number of unlabeled videos can be engaged in pseudo-labeling, facilitating each other’s training accordingly. Fig. 2 suggests that our CMPL obtains more high-quality pseudo-labels than the baseline. In addition, we also study the compatibility of this cross-model framework with conventional temporal data augmentations (*i.e.*, adjusting the temporal location and the frame rate), which are widely used in recent literature [5, 16, 29]. Experiments on a range of standard benchmarks and training settings demonstrate the effectiveness of our CMPL. In particular, when using only the RGB modality and 1% labeled videos, CMPL achieves

17.6% and 25.1% Top-1 accuracy on Kinetics-400 [2] and UCF-101 [18], surpassing FixMatch [17], by 9.0% and 10.3%, respectively. We also conduct a comprehensive empirical analysis to study how the cross-model supervision helps improve performance. This analysis shows that the primary backbone has large improvement on classes for which the auxiliary network works very well, supporting our motivation that the auxiliary network can complement the backbone.

2. Related Work

Semi-Supervised Learning for Image Classification. Semi-supervised learning has been well-explored in the area of image classification. Some prior works leverage consistency regularization which requires models to be robust to perturbations including data augmentations [25] and adversarial perturbations [13]. Recent work focuses more attention on the pseudo-labeling framework, which assigns labels to unlabeled images according to model predictions. In particular, [19] and [11] produce pseudo-labels by utilizing the exponential moving average of model parameters and historical predictions, respectively. Differently, FixMatch [17] combines consistency regularization and pseudo-labeling by requiring the predictions from strongly-augmented data to mirror those from weakly-augmented data. These methods do not explicitly consider the temporal dynamics that characterize human actions, which are shown in our work to be more effectively represented with the help of a separate complementary network.

Semi-Supervised Learning for Action Recognition. While action recognition [2, 4, 6, 7, 12, 14, 15, 20–24, 27, 30, 33, 34] has progressed significantly in recent years, semi-supervised learning for action recognition has been less studied. [8] represents actions by a conventional action bank and applies variants of extreme learning machines for final class prediction. [31] presents an encoder-decoder structure that is trained via an image reconstruction pretext task. [9] introduces a new framework that leverages a 2D image classifier to assist action recognition. [16] proposes a temporal contrastive learning framework to model temporal aspects by comparing the same video at different speeds.

A different approach is to learn from multiple views of the same source data. In the concurrent work of [28], the different views correspond to RGB, optical flow, and temporal gradients. Prediction results are separately generated from each view using a single common model, and these results are ensembled to produce pseudo-labels for retraining the model. This approach is similar to that of co-training [1], where two networks are asked to iteratively make predictions on the unlabeled data and the pseudo-labels are then merged to supervise the two models jointly. Differently, our approach requires the two models with different architectures to provide pseudo-labels for each

other in a cross-teaching manner. In this way, they are expected to benefit from the complementary representation learned by their counterparts.

3. Method

Given a labeled set $\mathcal{V} = \{(v_1, y_1), \dots, (v_{N_l}, y_{N_l})\}$ containing N_l videos and an unlabeled set $\mathcal{U} = \{u_1, \dots, u_{N_u}\}$ comprising N_u videos, CMPL learns an action recognition model by efficiently leveraging both labeled and unlabeled data. In general, N_u is much larger than N_l . In this section, we first present preliminary background on the pseudo-labeling framework in Sec. 3.1, then we introduce the proposed CMPL framework in Sec. 3.2, followed by a description of the implementation for CMPL in Sec. 3.3.

3.1. Preliminaries on Pseudo-Labeling

Pseudo-labeling is a commonly used technique for semi-supervised image recognition. It promotes the idea of utilizing the model itself to produce artificial labels for unlabeled data. The artificial labels obtained with confidence above a pre-defined threshold are retained, and the corresponding unlabeled data can then serve as additional training samples. A recent state-of-the-art method that adopts the pseudo-labeling scheme is FixMatch [17], which uses a weakly-augmented image as input to acquire the image’s pseudo-label, that is subsequently paired with a strongly-augmented version to form a labeled sample. FixMatch can be directly extended to semi-supervised action recognition as:

$$\mathcal{L}_u = \frac{1}{B_u} \sum_{i=1}^{B_u} \mathbb{1}(\max(p_i) \geq \tau) \mathcal{H}(\hat{y}_i, F(\mathcal{T}_{strong}(u_i))), \quad (1)$$

where B_u represents the batch size, τ is a threshold to determine whether a prediction is confident or not, $\mathbb{1}(\cdot)$ denotes the indicator function, $p_i = F(\mathcal{T}_{weak}(u_i))$ denotes class distribution, and $\hat{y}_i = \arg \max(p_i)$ represents the pseudo-label. $\mathcal{T}_{strong}(\cdot)$ and $\mathcal{T}_{weak}(\cdot)$ respectively denote the weak and the strong augmentation processes. $\mathcal{H}(\cdot, \cdot)$ denotes the standard cross-entropy loss. \mathcal{L}_u is the loss on the unlabeled data, and the loss on the labeled data is the cross-entropy loss commonly used in action recognition.

3.2. Cross-Model Pseudo-Labeling

As described in Sec. 3.1, the core idea of recent semi-supervised learning approaches is to construct as many high-quality pseudo-labels for unlabeled data as possible. However, when the number of labeled data is limited, the discriminative power of a single model is too weak to successfully assign a large number of pseudo labels to the unlabeled data. Therefore, our approach is to learn two models with different architectures and ask them to provide

pseudo-labels for each other, which is inspired by the observation that different models have different structural biases that lead to complementary semantic representations. As shown in Fig. 3, our CMPL framework employs two models (*i.e.*, the primary backbone $F(\cdot)$ and the auxiliary network $A(\cdot)$) with different capacities. They both learn from labeled data through supervised training and simultaneously provide their companion with pseudo-labels for unlabeled data. This symmetric design incentivizes the two models to learn complementary representations, which in turn boosts performance.

Learning on Labeled Data. It is straightforward to train a model on labeled data. Given a set of labeled videos $\{(v_i, y_i)\}_{i=1}^{B_l}$, we formulate the supervised loss of the two networks as:

$$\mathcal{L}_s^F = \frac{1}{B_l} \sum_{i=1}^{B_l} \mathcal{H}(y_i, F(\mathcal{T}_{standard}^F(v_i))), \quad (2)$$

$$\mathcal{L}_s^A = \frac{1}{B_l} \sum_{i=1}^{B_l} \mathcal{H}(y_i, A(\mathcal{T}_{standard}^A(v_i))), \quad (3)$$

where $\mathcal{T}_{standard}(\cdot)$ represents the standard data augmentation used in action recognition [4, 24].

Learning on Unlabeled Data. Given an unlabeled video u_i , the auxiliary network $A(\cdot)$ will make a prediction on the weakly-augmented u_i and output the category-wise probabilities, $p_i^A = A(\mathcal{T}_{weak}(u_i))$. If the maximum entry among these probabilities, $\max(p_i^A)$, surpasses a pre-defined threshold τ , we then regard this to be a solid prediction and utilize p_i^A to derive the pseudo ground-truth $\hat{y}_i^A = \arg \max(p_i^A)$ for the strongly-augmented u_i . In this way, the backbone $F(\cdot)$ can be learned with

$$\mathcal{L}_u^F = \frac{1}{B_u} \sum_{i=1}^{B_u} \mathbb{1}(\max(p_i^A) \geq \tau) \mathcal{H}(\hat{y}_i^A, F(\mathcal{T}_{strong}(u_i))), \quad (4)$$

where B_u denotes batch size, and $\mathcal{H}(\cdot, \cdot)$ is cross-entropy loss.

Similarly, the primary backbone will also make a prediction $p_i^F = F(\mathcal{T}_{weak}(u_i))$ and use it to generate a labeled pair $(\hat{y}_i^F, \mathcal{T}_{strong}(u_i))$ for the auxiliary network:

$$\mathcal{L}_u^A = \frac{1}{B_u} \sum_{i=1}^{B_u} \mathbb{1}(\max(p_i^F) \geq \tau) \mathcal{H}(\hat{y}_i^F, A(\mathcal{T}_{strong}(u_i))). \quad (5)$$

Complete Training Objective. To summarize, with the supervised losses from labeled data and the unsupervised losses from unlabeled data, the complete objective function is presented as

$$\mathcal{L} = (\mathcal{L}_s^F + \mathcal{L}_s^A) + \lambda(\mathcal{L}_u^F + \mathcal{L}_u^A), \quad (6)$$

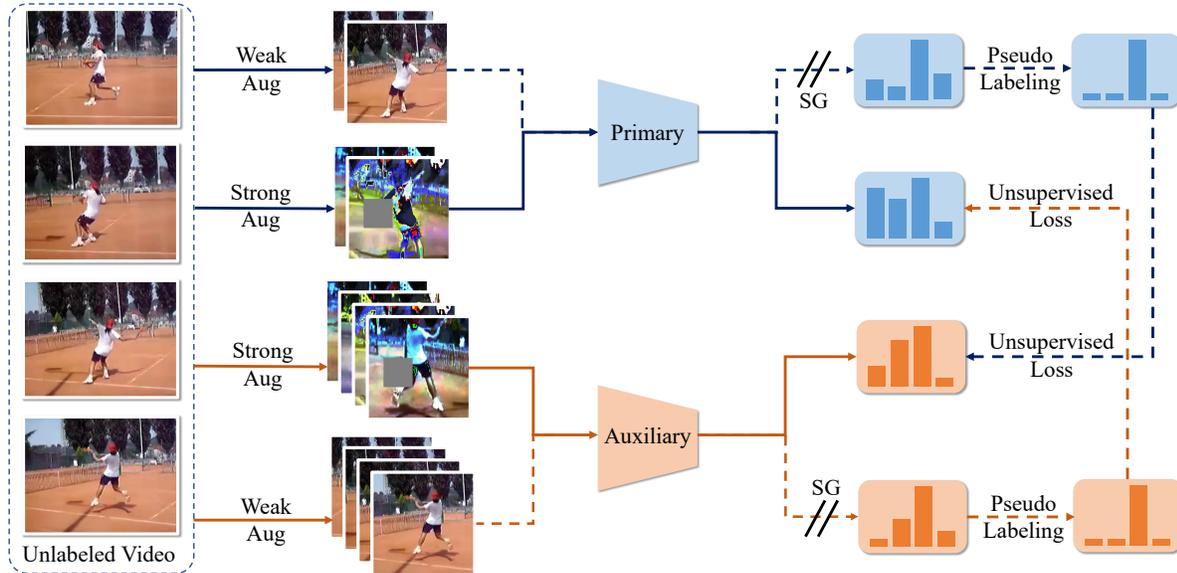


Figure 3. **Illustration of the proposed Cross-Model Pseudo-Labeling (CMPL) framework**, which consists of two models, *i.e.*, the primary backbone $F(\cdot)$ and the auxiliary network $A(\cdot)$. They take video input at different frame rates. Unless specified, we adopt 3D-ResNet50 as the primary backbone and a lightweight 3D-ResNet50 $\times 1/4$ as the auxiliary network. Given an unlabeled video, these two models make independent predictions on weakly-augmented data. The predicted results are used to produce a pseudo-label for their sibling where it is used as supervision for a strongly-augmented version. “SG” stands for the stop-gradient operation. The supervised losses from labeled data are omitted in this figure.

where λ denotes the balancing coefficient for the unsupervised losses. Note that the size of the auxiliary network $A(\cdot)$ is designed to be much smaller than the backbone $F(\cdot)$; therefore, the introduction of $A(\cdot)$ has little effect on training efficiency.

3.3. Implementation

In this part, we describe the implementation of CMPL.

Auxiliary Network. As stated in Sec. 3.2, the auxiliary network should have a different capacity from the primary network to provide complementary representations. In practice, we obtain the auxiliary network by adjusting the depth and width of the primary network. For instance, if the primary network is 3D-ResNet50, the auxiliary network could be 3D-ResNet18 or 3D-ResNet50 $\times 1/4$ where the channels are reduced four-fold. Comprehensive ablation studies on the depth and width of the auxiliary network are included in Tab. 3 of the experiments section. Unless otherwise specified, we adopt 3D-ResNet50 and 3D-ResNet50 $\times 1/4$ as the primary backbone and auxiliary network, respectively.

Temporal Data Augmentations. For spatial data augmentation, we strictly follow the augmentations of [4, 24] for training as $\mathcal{T}_{standard}(\cdot)$ on labeled data. For unlabeled data, center cropping with a patch size of 224×224 is employed as the weak augmentation $\mathcal{T}_{weak}(\cdot)$. Following FixMatch [17], RandomAugmentation [17, 25] together with Cutout

[3] are used as the strong augmentations $\mathcal{T}_{strong}(\cdot)$. Note that all the spatial transformations of the above augmentations are temporally consistent across all frames in the same video clip.

Besides spatial augmentations, commonly used temporal data augmentations, namely various *temporal locations* and *frame rates*, are compatible with our CMPL. Previous work [5] suggests that visual content is often temporally persistent throughout the time span of a video, such that different clips of a given video share similar representations. Accordingly, we adjust the temporal location of the input clip and control the augmentation strength via an adjustable **time offset** t_s . In our CMPL framework, this is done by feeding the primary backbone $F(\cdot)$ and the auxiliary network $A(\cdot)$ clips from different temporal locations of the same video while still requiring them to supervise each other. Meanwhile, we also follow [16, 29] in regarding different frame rates as a form of temporal augmentation. This is also illustrated in Fig. 3, where a faster frame rate results in more frames given that all clips share the same temporal duration.

4. Experiments

We evaluate the proposed CMPL on two commonly used datasets, namely Kinetics-400 [2] and UCF-101 [18], using two standard settings for semi-supervised action recognition (*i.e.* 1% and 10% labeled data). Detailed

Table 1. **Comparison with state-of-the-art methods on UCF-101 and Kinetics-400.** Note that R18 and R50 denote the backbone networks and their depths. We report top-1 accuracy as the evaluation metric. “Input” shows the input data format used during training, where “V” is the raw rgb video, “F” is optical flow and “G” is the temporal gradient.

Method	Backbone	Input	#Frames	UCF-101		Kinetics-400	
				1%	10%	1%	10%
Supervised	3D-ResNet50	V	8	6.5	32.4	4.4	36.2
<i>Image-based Semi-supervised Methods</i>							
S4L [32]	3D-ResNet18	V	16	-	29.1	-	-
FixMatch [17]	3D-ResNet50	V	8	14.8	49.8	8.6	46.9
FixMatch [17]	SlowFast-R50	V	8	16.1	55.1	10.1	49.4
<i>Video-based Semi-supervised Methods</i>							
VideoSSL [10]	3D-ResNet18	V	16	-	42.0	-	-
TCL [16]	TSM-R18	R	8	-	-	8.5	-
MvPL [28]	3D-ResNet50	V+F+G	8	22.8	80.5	17.0	58.2
Ours	3D-ResNet18	V	8	23.8	67.6	16.5	53.7
Ours	3D-ResNet50	V	8	25.1	79.1	17.6	58.4

ablation studies on the design choices of CMPL are also included. In addition, some empirical analysis is presented to validate our motivations for CMPL. We note that all the experiments are conducted on a single modality (*i.e.*, RGB) and evaluated on the corresponding validation set unless otherwise stated.

4.1. Settings

Dataset. Kinetics-400 [2] consists of 240K and 20K videos respectively for training and validation, covering 400 categories. Each video lasts about 10 seconds. UCF-101 [18] contains 101 action classes with roughly 9.5K training videos and 4K validation videos. To obtain 1% or 10% labeled data for the two evaluation settings, for Kinetics-400 we form two balanced labeled subsets by sampling 6 or 60 videos per class. As for UCF101, we sample 1 or 10 videos per class as labeled data. To reduce the randomness brought by the sampling process, the accuracy averaged over **three different sampled subsets** is reported. To obtain the input frames for each video, we first extract 64 consecutive frames as a raw clip, then sample sparsely with a stride s to obtain $64/s$ frames, denoted as $64/s \times s$. Unless stated otherwise, the input frames for the primary and auxiliary branches are set to 8×8 and 16×4 , respectively.

Training. For Kinetics-400, we train the model using 16 GPUs, with the SGD optimizer, momentum 0.9, and weight decay 0.0001. The batch size for labeled data and unlabeled data are 2 and 10 on each GPU, resulting in a mini-batch size of 192 in total. The base learning rate is set to 0.1 and decayed according to the cosine scheduler. A total of 200 training epochs is used. In addition, λ is set to 5 and 2 for the 1% and 10% labeled data settings, and τ is set to 0.9. The training setting of UCF-101 is the same as that for

Kinetics-400 except that the weight decay is set to 0.001.

Inference. Following the test protocol in [4, 24, 30], we uniformly sample ten clips over the whole video and average the softmax probabilities of all clips as the final prediction. The shorter spatial side of the input video is scaled to 256 pixels, and three crops of size 256×256 are extracted to obtain more spatial information. Although the primary and auxiliary network are jointly optimized in training, we use only the primary model for inference, leading to no additional inference cost.

Baselines. First, we have a supervised baseline with the same architecture as our approach and the same set of labeled samples for training, but without any unlabeled data. Secondly, we extend the state-of-the-art semi-supervised learning approaches in the image domain [17, 32] to the video domain for comparison. For Fixmatch [17], we adopt the same experimental settings as our approach. Finally, we also include the state-of-the-art video-based semi-supervised learning methods [10, 16, 28], whose performances from their original papers are directly reported.

4.2. Main Results

Quantitative results are presented in Tab. 1. Compared to previous methods that use a single modality of RGB frames (all baselines except for MvPL [28]), our CMPL leads by a clear margin even when equipped with a shallower network, *i.e.*, 3D-ResNet18 as the backbone. Further improvements are obtained by increasing the capacity of the backbone network *i.e.*, from 3D-ResNet18 to 3D-ResNet50.

While CMPL outperforms FixMatch when both of these methods adopt 3D-ResNet50 as the backbone, the improvement of CMPL may come from the inclusion of the lightweight auxiliary network. To investigate this, we

Table 2. **Ablation studies on CMPL components.** We gradually add the auxiliary network (Auxiliary), different visual tempos (Frame Rates) and different temporal locations (Temporal Loc) to the baseline model (FixMatch). All studies are trained with **50 epochs on 1% labeled Kinetics-400**. Unless specified, the results of following tables are all with the same training configuration.

Auxiliary	Frame Rates	Temporal Loc	Top-1
	<i>FixMatch</i>		6.78
	✓		7.12
	✓	✓	7.68
✓			12.04
✓	✓		12.90
✓	✓	✓	13.71

change the backbone of FixMatch from 3D-ResNet50 to SlowFast-R50, which is a two-branch network that also includes a lightweight component for feature aggregation. As shown in Tab. 1, CMPL still outperforms FixMatch with a larger backbone, suggesting it is the auxiliary supervision that matters.

Meanwhile, the concurrent work MvPL [28] achieves performance competitive to ours. However, for these results it uses 3 modalities, compared to only the RGB modality for CMPL.

4.3. Ablation Study

In this part, we make comprehensive analysis of the proposed CMPL via ablation studies. We first analyze the key components of our approach in Sec. 4.3.1, followed by the experimental results of different instantiations of losses and hyperparameters in Sec. 4.3.2. Note that all ablation studies are performed on **1% labeled data of Kinetics-400 with 50 training epochs**. The 3D-ResNet50 backbone and the sparse sampling strategy (8×8) are adopted for the primary backbone unless specified otherwise.

4.3.1 CMPL Components

Ablations on CMPL components. While the auxiliary network provides pseudo-label estimates that complement those generated by the primary backbone, we adopt two advanced temporal augmentations to further enhance divergence between the representations of these two networks, to encourage greater complementarity. As shown in Tab. 2, starting from the baseline approach FixMatch which achieves 6.78% Top-1 accuracy, introducing the auxiliary supervision improves the accuracy to 12.04%, demonstrating the effectiveness of the cross-model strategy. On top of this, the temporal augmentations further enhance performance to 13.71%, leading to a stronger result.

Auxiliary network architecture. To study the effect of auxiliary network architecture, we start by using the

Table 3. **Architectures of the auxiliary network.** Different architectures are adopted as the auxiliary network to study the effects of network architectures. Specifically, when the auxiliary pathway adopts the primary’s architecture, CMPL degenerates to FixMatch [17], providing a strong baseline for our comparison.

Auxiliary Network	#Frames	#FLOPs	Top-1
Primary	8	146G	6.78
Primary	16	168G	7.12
3D-ResNet50	16	187G	9.97
<i>network depth</i>			
3D-ResNet18	16	132G	10.63
3D-ResNet10	16	108G	10.88
<i>network width</i>			
3D-ResNet50 \times 1/2	16	112G	11.57
3D-ResNet50 \times 1/4	16	94G	12.90
3D-ResNet50 \times 1/8	16	87G	12.08

Table 4. **Effects of input #Frames.** Different number of frames are adopted as the auxiliary input to evaluate the effect of frame rates in CMPL.

T	8×8	16×4	32×2
Top-1	12.04	12.90	12.94

same architecture for both the backbone and the auxiliary network, which leads to two different versions, one with a shared set of parameters (equivalent to FixMatch) and another with two different sets of parameters, corresponding to the 2nd and 3rd lines of Tab. 3.

Subsequently, we also shrink the depth and width of 3D-ResNet50 to produce other auxiliary network variants. Tab. 3 presents the performance with different auxiliary networks. It can be seen that even with the backbone and auxiliary network being identical but with different weights, learning in the cross-model manner still leads to performance improvements. Moreover, substantial gains are obtained no matter which factor (*e.g.*, depth or width) is shrunk, indicating that a relatively small network as the auxiliary network can provide complementary information for the backbone. In addition, the cross-model learning strategy is the key to the enhanced performance, rather than the newly introduced model capacity, since consistent gains are observed as the auxiliary network varies.

Frame rate for visual tempo. To study how the frame rates between the primary backbone and auxiliary network affect the performance, we fix the sampling rate (8×8) for the backbone, and use different sampling rates for the auxiliary network, namely 8×8 , 16×4 , and 32×2 .

As shown in Tab. 4, a larger sampling rate leads to a better performance. As stated in previous work [4, 30], a larger difference between two sampling rates usually leads

Table 5. **Effects of time offset t_s .** The time offset determines the temporal distance between the clips processed by the primary and auxiliary networks.

t_s	1	2	5	10
Top-1	12.90	12.64	13.44	13.71

Table 6. **Effects of the losses for the auxiliary network.**

\mathcal{L}_s^A	\mathcal{L}_u^A	Top-1	Top-5
	<i>FixMatch</i>	6.78	17.92
✓		10.57	25.53
	✓	10.32	24.46
✓	✓	12.90	28.19

to a more different representations of the same videos at the input level. This naturally constructs two relevant and complementary views, further boosting our cross-model learning.

Time offset for temporal location. Besides frame rate, we also explore the effect of the time offset for temporal location, which is controlled by the hyper-parameter t_s as shown in Sec. 3.3. Tab. 5 reports the performance for increasing t_s . It is seen that a larger t_s positively affects the final performance. It should be noted that the trimmed videos of Kinetics last about 10s so that t_s is set to at most 10.

4.3.2 Choices of losses and hyperparameters

Losses for auxiliary branch. There is a supervised and an unsupervised loss for both the primary and the auxiliary networks. Here, we examine the importance of the supervised and unsupervised loss of the auxiliary network, while the two losses of the primary network are kept.

Tab. 6 presents the results. Including the auxiliary network always improves performance, no matter where the labels come from (e.g. pseudo-labels or ground-truth labels). This demonstrates that the auxiliary network brings representations complementary to those of the primary network. Including both types of labels leads to the best results.

Hyperparameters. Fig. 4a presents the effects of different thresholds, where using 0.9 as the threshold obtains the best performance. It is concluded that a higher threshold leads to better performance, i.e., the quality of the pseudo labels matters. Besides, we also consider the ratio of unlabeled data to labeled data (B_u/B_l). We fix the batch size of labeled data to be 2 and set the ratio to be $\{1, 2, 5, 10\}$ to analyze its effect. As shown in Fig. 4b, B_u/B_l set to be 5 can achieve the best trade-off between performance and computation cost. We also analyze the effect of the loss weight λ and observe that it affects the performance as

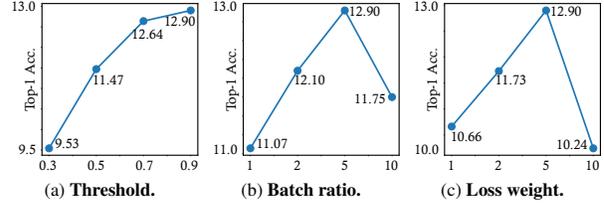


Figure 4. **Effects of hyperparameters.** Results of varying threshold, the ratio of unlabeled data to labeled data (B_u/B_l), as well as the loss weight (μ), are included to comprehensively study the effects of the hyperparameters.

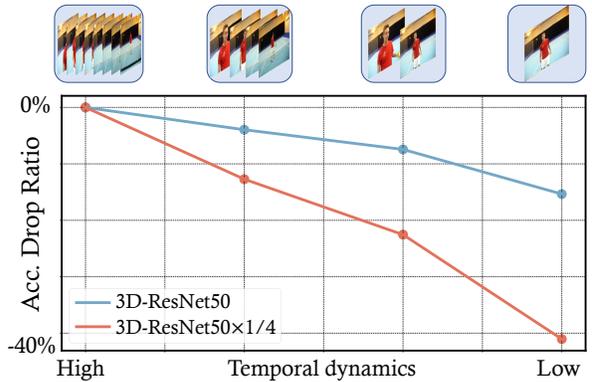


Figure 5. **Accuracy drop ratio with different frames per clip.** The top half of this figure presents the re-sampled frames with different strides. As the stride increases, the number of re-sampled frames decreases, weakening the temporal dynamics of the given clip. We evaluate the performance drop ratio of the re-sampled frames to study the representations learned by the two different networks.

accuracy drops by 2% when setting λ to 1 or 10. Fig. 4c illustrates that the optimal value of the loss weight is 5, at which the supervised loss and unsupervised loss is properly balanced.

4.4. Empirical Analysis

To validate our motivations for CMPL, some empirical analysis is conducted on Kinetics-400 with 1% labeled data. Unless specified otherwise, we use 3D-ResNet50 as the backbone, which takes 8 frames as input. As for the auxiliary network, we use 3D-ResNet50x1/4, which takes 16 frames as input.

Representation learned by different networks. Fig. 1 indicates that networks of different scale learn complementary semantics with limited annotated data. To study the mechanism by which these two networks complement each other, we first train the primary and auxiliary networks separately with 8×8 frames as the input. They are only trained with 1% labeled samples of Kinetics-400. We then re-sample the 8 frames with different stride $\{1, 2, 4, 8\}$ in order to decrease the temporal dynamics of a given video clip. Notably, no matter how many frames the sampled

clip contains, they are always extended to a 8-frame clip for testing.

Fig. 5 displays accuracy curves of the two networks with inputs of varying temporal dynamics. It can be seen that the auxiliary network is more sensitive to temporal dynamics. This suggests that the tiny model using a lower channel capacity can better capture fast motion without building a detailed spatial representation, providing complementary temporal semantics in relation to the spatial patterns learned by the backbone network.

Effects of the auxiliary pseudo-labels. A key difference between CMPL and FixMatch [17] is the source of the pseudo-labels. The primary backbone takes the pseudo-labels from the auxiliary network while FixMatch only uses its own high-confidence predictions. To study the quality of the pseudo-labels provided by the auxiliary network, we compare three sources of pseudo-labels, namely the backbone network, the auxiliary network trained in CMPL, and the 3D-ResNet50 trained in FixMatch, which has the same architecture as the backbone network. We first select the samples with pseudo-labels assigned by the auxiliary network to build a subset. And then the accuracy of these three pseudo-label sources are tested on this subset. Accuracy is recorded over the whole training process at an interval of 10 epochs.

Fig. 6 shows the accuracy curves. Obviously, for the samples originally labeled by the auxiliary network with high confidence, the auxiliary network consistently obtains better pseudo-labels (orange curve) than the backbone network (blue curve) throughout the training process. Moreover, compared to the same architecture learned with FixMatch (green curve), the backbone network gradually produces better estimates for these selected samples. Such results demonstrate 1) cross-model pseudo-labeling using a tiny auxiliary network can enhance the backbone network; 2) the backbone network indeed learns from the auxiliary network through the proposed cross-model strategy; 3) although the low-capacity auxiliary network tends to perform worse than the high-capacity backbone network, it still provides knowledge to guide the latter.

Primary gain vs. Auxiliary accuracy. The key idea of CMPL is to adopt an auxiliary network to provide complementary information that promotes better representation. Therefore, we examine the per-class gain (*primary gain*) of the backbone network on the Kinetics-400 validation dataset when adopting our CMPL instead of FixMatch. We also include the per-class performance of the auxiliary network (*auxiliary accuracy*) trained in a supervised manner as a reference. For better visualization, we discretize auxiliary accuracy into bins with an interval of 0.05 and then calculate the mean of the primary gain (y-axis) in each bin. As shown in Fig. 7, the performance gain of the backbone network is positively correlated with that of

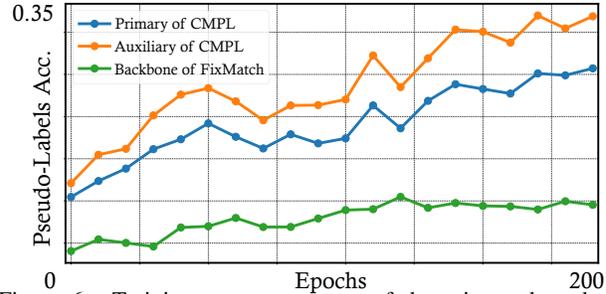


Figure 6. Training accuracy curves of the primary branch of CMPL (blue), the backbone of FixMatch (green), and the auxiliary of CMPL (orange). Training accuracy is evaluated on the samples with pseudo-labels assigned by the auxiliary network. Please refer to Sec. 4.4 for more details.

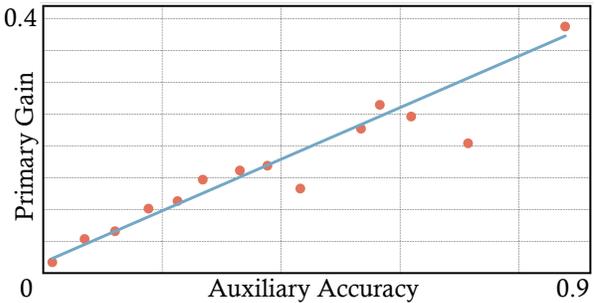


Figure 7. **Primary gain vs. Auxiliary accuracy.** Each red point denotes the mean accuracy gain of the primary backbone within a bin of auxiliary network accuracy. The blue line plots a least squares approximation.

the auxiliary network on the corresponding classes. This study supports our motivation that the auxiliary network complements the backbone network, particularly on action classes recognized well by the former.

5. Conclusion

We present a Cross-Model Pseudo-Labeling framework for semi-supervised action recognition. It pairs a lightweight auxiliary network with the primary backbone such that they can learn from each other via pseudo-labeling. Over multiple datasets and low data regimes, our approach outperforms its supervised counterpart with limited labeled data by a substantial margin and surpasses several previous methods using only RGB frames. Moreover, we show that the primary network can be complemented by an auxiliary network even with the same input, demonstrating a new alternative to the previous standard of changing the input data in semi-supervised learning. While CMPL obtains improved performance for semi-supervised action recognition, its results with 10% annotated data are not as good as supervised training with 100% labeled data. This leaves room for further investigation into improving CMPL with more advanced techniques.

Acknowledgement We thank Nanxuan Zhao for discussion and comments about this work.

References

- [1] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, 1998.
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017.
- [3] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with dropout. *arXiv preprint arXiv:1708.04552*, 2017.
- [4] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *Int. Conf. Comput. Vis.*, 2019.
- [5] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3299–3309, 2021.
- [6] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal multiplier networks for video action recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017.
- [7] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016.
- [8] Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas. Semi-supervised classification of human actions based on neural networks. In *Int. Conf. Pattern Recog. IEEE*, 2014.
- [9] Longlong Jing, Toufiq Parag, Zhe Wu, Yingli Tian, and Hongcheng Wang. Videoss1: Semi-supervised learning for video classification. In *IEEE Winter Conf. Appl. Comput. Vis.*, 2021.
- [10] Longlong Jing, Toufiq Parag, Zhe Wu, Yingli Tian, and Hongcheng Wang. Videoss1: Semi-supervised learning for video classification. In *IEEE Winter Conf. Appl. Comput. Vis.*, 2021.
- [11] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Int. Conf. Mach. Learn. Worksh.*, 2013.
- [12] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Int. Conf. Comput. Vis.*, 2019.
- [13] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018.
- [14] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatiotemporal representation with pseudo-3d residual networks. In *Int. Conf. Comput. Vis.*, 2017.
- [15] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Adv. Neural Inform. Process. Syst.*, 2014.
- [16] Ankit Singh, Omprakash Chakraborty, Ashutosh Varshney, Rameswar Panda, Rogerio Feris, Kate Saenko, and Abir Das. Semi-supervised action recognition with temporal contrastive learning. *arXiv preprint arXiv:2102.02751*, 2021.
- [17] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.
- [18] 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.
- [19] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Adv. Neural Inform. Process. Syst.*, 2017.
- [20] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Int. Conf. Comput. Vis.*, 2015.
- [21] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *Int. Conf. Comput. Vis.*, 2013.
- [22] Limin Wang, Yu Qiao, and Xiaoou Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015.
- [23] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Eur. Conf. Comput. Vis.*, 2016.
- [24] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018.
- [25] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*, 2019.
- [26] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [27] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Eur. Conf. Comput. Vis.*, 2018.
- [28] Bo Xiong, Haoqi Fan, Kristen Grauman, and Christoph Feichtenhofer. Multiview pseudo-labeling for semi-supervised learning from video. *arXiv preprint arXiv:2104.00682*, 2021.
- [29] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video representation learning with visual tempo consistency. *arXiv preprint arXiv:2006.15489*, 2020.
- [30] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [31] Ming Zeng, Tong Yu, Xiao Wang, Le T Nguyen, Ole J Mengshoel, and Ian Lane. Semi-supervised convolutional neural networks for human activity recognition. In *IEEE International Conference on Big Data*. IEEE, 2017.
- [32] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1476–1485, 2019.

- [33] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Trajectory convolution for action recognition. In *Adv. Neural Inform. Process. Syst.*, 2018.
- [34] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *Eur. Conf. Comput. Vis.*, 2018.