

Attention Distillation for Learning Video Representations

Miao Liu¹

mliu328@gatech.edu

Xin Chen³

xinchen@kwai.com

Yun Zhang¹

yzhang467@gatech.edu

Yin Li²

yin.li@wisc.edu

James M. Rehg¹

rehg@gatech.edu

¹ College of Computing

Georgia Institute of Technology
Atlanta, United States

² Departments of Biostatistics
and of Computer Sciences

University of Wisconsin-Madison
Madison, United States

³ Y-tech

Kuaishou Technology
Palo Alto, United States

Abstract

We address the challenging problem of learning motion representations using deep models for video recognition. To this end, we make use of attention modules that learn to highlight regions in the video and aggregate features for recognition. Specifically, we propose to leverage output attention maps as a vehicle to transfer the learned representation from a flow network to an RGB network. We systematically study the design of attention modules, develop a novel method for attention distillation, and evaluate our method on major action recognition benchmarks. Our results suggest that our method improves the performance of the baseline RGB network by a significant margin while maintains similar efficiency. Moreover, we demonstrate that attention serves a more robust tool for knowledge distillation in video domain. We believe our method provides a step forward towards learning motion-aware representations in deep models and valuable insights for knowledge distillation. Our project page is available at <https://aptx4869lm.github.io/AttentionDistillation/>

1 Introduction

Action recognition in videos has emerged as a key challenge for deep models. This task requires the understanding of both spatial and temporal cues and the best methods for extracting and fusing them. The two-stream architecture [5], exemplified by the I3D model [9], has proven to be an effective framework for addressing these challenges. Fusing two modalities of appearance and motion is conceptually appealing, yet it is computationally expensive. A two-stream model can be 100 times slower than its single RGB stream version [6]. Moreover, learning motion-aware video features from RGB frames remains a challenging problem [8]. In this context, we address the following research questions: *Does a deep model need an explicit flow channel to capture motion patterns? How can we bridge the gap between an*

RGB stream network and its two stream version without incurring the extra computational cost? Several previous works have addressed the challenge of learning a video representation that encodes motion information using a single RGB stream [5, 19, 39, 40, 42]. Our work shares the same motivation, but pursues a very different approach.

We present a novel video representation learning method called *attention distillation*. Our method makes use of an explicit probabilistic attention model, and leverages motion information available at training time to predict the motion-sensitive attention features from a single RGB stream. In addition to their utility in visualizing and understanding learned feature representations, we argue that attention models provide an attractive vehicle for mapping between sensing modalities in a task-sensitive way. Once learned, our model requires only RGB frames as inference inputs, and jointly predicts appearance and motion attention maps for action recognition. We conduct extensive experiments and demonstrate that our attention distillation enables more accurate action recognition across several video datasets, while remaining very efficient.

Our main contributions are summarized as follows:

- We propose a novel method for learning motion-aware video representations from RGB frames. Our method distills motion knowledge into an RGB network by mimicking the attention map of a reference flow network.
- Our method achieves consistent improvements of $\sim 1\%$ across major datasets, including UCF101 [68], HMDB51 [21], EGTEA [22], and 20BN-V2 [29], with almost no extra computational cost.
- We study different choices of attention modules for action recognition, and demonstrate that attention serves as a more robust vehicle for knowledge distillation in comparison to previous feature distillation methods.

2 Related Works

2.1 Action Recognition

Action recognition is well studied in computer vision [34]. Recent efforts focus on developing novel deep models for action recognition. For example, recent works [3, 14, 41] proposed to make use of 3D convolutional networks to capture spatio-temporal features beyond a single frame. However, their performance using video frames alone falls far behind their two stream versions [9]. Our work seeks to address the problem of recovering the motion cues encoded in videos from RGB frames alone. There are several recent attempts in this direction. Bilen et al. [1] proposed a dynamic image network that makes use of the parameters of a ranking machine that captures the temporal evolution of the video frames. Ng et al. [63] proposed to jointly predict action labels and flow maps from video frames using multi-task learning. This idea is extended by Fan et al. [12], where they fold the TV-L1 flow estimation [32] into their TVNet. Without using flow, Tran et al. [42] demonstrated that factorized 3D convolutions (2D spatial convolution and 1D temporal convolution) can facilitate the learning of spatio-temporal features. A similar finding was also presented by Xie et al. [46]. Our method shares the same motivation as these approaches, yet takes a vastly different route. We explore attention mechanisms for video recognition, and propose to distill the predicted attention from a flow network to an RGB network.

2.2 Knowledge Distillation

Our attention distillation method is inspired by knowledge distillation, first proposed by [2] for model compression and further popularized by [16]. Some recent works [10, 13, 27] explored knowledge distillation across modalities. The most relevant works are [8, 39]. They both addressed the challenge of video representation learning via knowledge distillation. They assume that the reference flow network has better performance than the RGB stream network, and seek to regularize the learning of the RGB stream by distilling features from the flow network to the RGB stream. However, for recently developed large scale video datasets (e.g., Kinetics [20], Charades [36]) and egocentric video datasets (e.g., EGTEA [2] and EPIC-Kitchens [6]), the RGB stream has better performance than the flow stream. Feature distillation methods also suffer from the potential threat of “overwriting” the features from RGB stream with the features from flow stream. This is related to the pitfall of catastrophic forgetting [9]. In contrast, we propose to distill attention maps— which are indicators of important regions for recognition. This design choice stems from the key challenge of video representation learning—motion is substantially different from appearance and both modalities are important for recognition. Our experimental results in Sec. 4 demonstrate that our method can overcome the disadvantages of previous feature distillation methods.

2.3 Attention for Recognition

Attention has been widely used for visual recognition. We focus on selective visual attention that highlights discriminative regions. This is very different from the recent efforts on self-attention, i.e., self-similarity [12, 14, 49]. Recently, selective attention has been explored in deep models for object recognition [50] and image captioning [47]. Attention enables these models to “fixate” on image regions, where the decision is made based on a sequence of fixations. Several attention mechanisms are proposed for deep models. For example, [35] integrated soft attention in LSTMs for action recognition. [24] further extends [47] into videos. Specifically, they combined LSTMs with motion-based attention to infer the location of the actions. [11] modeled top-down and bottom-up attention using bilinear pooling. [44] proposed a residual architecture for soft attentions. Our previous work [22, 23] considered attention as a probabilistic distribution for egocentric action recognition. Our recent work [25] made use of motor attention for action anticipation. In this paper, we demonstrate that a useful probabilistic attention model can be obtained without access to a prior distribution from human gaze data. We also provide a systematical study of the utility of probabilistic attention model in action recognition.

3 Distilling Motion Attention for Actions

In this section, we present our method of attention distillation. We start with an overview of the key ideas, followed by a detailed description of the components in our method. Finally, we describe our network architecture and discuss the implementation details.

3.1 Overview

For simplicity, we consider an input video with a fixed length of T frames. Our method can easily generalize to multiple videos, e.g., for mini-batch training. We denote the input video

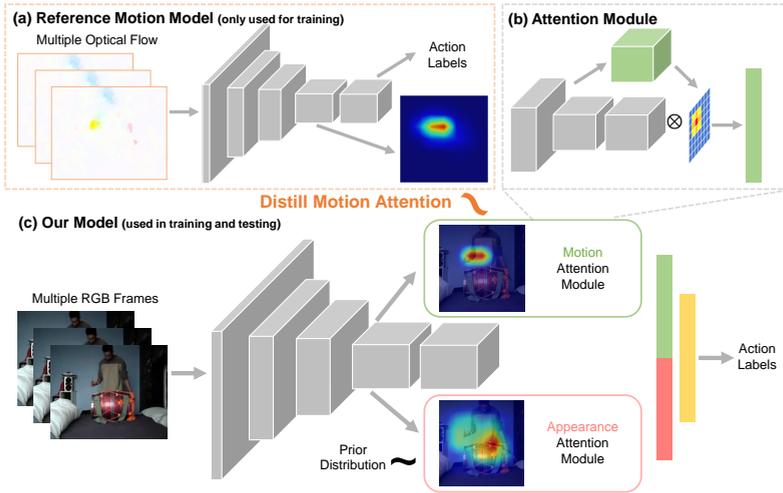


Figure 1: Overview of our method. Our model (c) takes multiple RGB frames as inputs and adopts a 3D convolutional network as the backbone. It outputs two attention maps using the attention module (b), based on which the action labels are predicted. The motion map is learned by mimicking the attention from a reference flow network (a). The appearance map is learned to highlight discriminative regions for recognition. These two maps are used to create spatio-temporal feature representations from video frames for action recognition.

as $x = \{x^1, x^2, \dots, x^T\}$, where x^t is a frame of resolution $H \times W$ with t as the frame number. Given x , our goal is to predict a video-level action label y . We leverage the intermediate output of a 3D convolutional network ϕ to represent x . This is given by a 4D tensor $\phi(x)$ of the size $T_\phi \times H_\phi \times W_\phi \times C_\phi$. C_ϕ is the feature dimension of 3D grids $T_\phi \times H_\phi \times W_\phi$ from the video x . Our method consists of three key components:

- **Attention Generation.** Our model first predicts an attention map \mathcal{A} based on $\phi(x)$ using the attention mapping function F_A . \mathcal{A} is a 3D tensor of size $T_\phi \times H_\phi \times W_\phi$. Moreover, \mathcal{A} is normalized within each temporal slice, i.e., $\sum_{w,h} \mathcal{A}(t, w, h) = 1$. \mathcal{A} is thus a sequence of 2D attention maps $\mathcal{A}(t)$ defined over T_ϕ steps.

- **Attention Guided Recognition.** Based on the attention map \mathcal{A} and the feature map $\phi(x)$, our model further applies a recognition module F_R to predict the action label y . Specifically, this module uses \mathcal{A} to selectively pool features from $\phi(x)$, followed by a classifier that maps the result feature vectors to the action label y .

- **Attention Distillation.** To regularize the learning, we assume that \mathcal{A} will receive supervision from a teacher model that outputs a reference attention map $\tilde{\mathcal{A}}$. The teacher model comes from the flow stream and is equipped with the same attention module for recognition.

Fig. 1 presents an overview of our method. Our model takes multiple video frames x as inputs, and learns to predict two attention maps based on $\phi(x)$: \mathcal{A}^M for motion attention and \mathcal{A}^A for appearance attention. Based on these two maps, the model further aggregates visual features that will be passed into the final recognition sub-network. During training, we match \mathcal{A}^M to the attention map $\tilde{\mathcal{A}}^M$ from the reference flow network. For testing, only the input video is required for recognition. Our model also outputs two attention maps that can be used to visualize and diagnose recognition performance. We now detail the design of our key components.

3.2 Attention Generation

We explore two different approaches for generating an attention map from the features $\phi(x)$, including soft attention [14] and its probabilistic version [2].

Soft Attention. Attention maps can be created by a linear function of $w_a \in R^{C_\phi}$ over the feature map $\phi(x)$,

$$F_{\mathcal{A}}(\phi(x)) = \text{softmax}(w_a * \phi(x)), \quad (1)$$

where $*$ is the 1x1 convolution on 3D feature grids. Softmax is applied on every time slice to normalize each 2D map.

Probabilistic Soft Attention. An alternative approach is to further model the distribution of linear mapping outputs as discussed in [2], namely

$$\mathcal{A} \sim p(\mathcal{A}) = \text{softmax}(w_a * \phi(x)) \quad (2)$$

where we model the distribution of \mathcal{A} . During training, an attention map can be sampled from $p(\mathcal{A})$ using Gumbel Softmax trick [18, 28]. We follow [2] to regularize the learning by adding additional loss term of

$$\mathcal{L}^R = \sum_t KL[\mathcal{A}(t) || U], \quad (3)$$

where $KL[\cdot]$ is the Kullback-Leibler divergence and U is the 2D uniform distribution ($H_\phi \times W_\phi$). This term matches each time slice of the attention map to the prior distribution. It is derived from variational learning and accounts for (1) the prior of attention maps and (2) additional regularization by spatial dropout [2]. During testing, we directly plug in $p(\mathcal{A})$ (the expected value of \mathcal{A}) for approximate inference.

Note that for both approaches, we restrict $F_{\mathcal{A}}$ to a linear mapping without a bias term. In practice, this linear mapping avoids the trivial solution of generating a uniform attention map by setting w to all zeros. This all-zero solution almost never arises during training when using a proper initialization of w .

3.3 Attention Guided Recognition

Our recognition module makes use of an attention map \mathcal{A} to select features from $\phi(x)$. Again, we consider two different models for the attention guided recognition.

Attention Pooling. Inspired by [27, 24], we design the function $F_{\mathcal{R}}$ as

$$\bar{y} = F_{\mathcal{R}}(\phi(x), \mathcal{A}) = \text{softmax}\left(W_r^T (\mathcal{A} \otimes \phi(x))\right) \quad (4)$$

where \otimes denotes the tilted multiplication $\mathcal{A} \otimes \phi(x) = \sum_{t,h,w} \mathcal{A}(t,h,w) \phi(x)_{t,h,w,c}$. This operation is equivalent to weighted average pooling with the weights shared across all channels.

Residual Connection. Using the attention map to re-weight features helps to filter out background noise, yet may also increase the potential risk of missing important foreground features. This drawback was discussed in [24]. We follow their solution of using a residual connection to the attention map, given by

$$\bar{y} = F_{\mathcal{R}}(\phi(x), \mathcal{A}) = \text{softmax}\left(W_r^T ((\mathcal{A} + I) \otimes \phi(x))\right), \quad (5)$$

where I is a 3D tensor of all ones. Intuitively, this operation further adds average pooled features to the representation before the linear classifier. By adding the residual term, the features learned by the network are preserved.

3.4 Attention Distillation

The key to our approach lies in the use of attention distillation during training. Specifically, we assume that a reference flow network is given as the teacher network. The teacher model also uses an attention mechanism for recognition. Moreover, its motion attention map $\tilde{\mathcal{A}}^M$ is used as additional supervisory signal for training our RGB network. This RGB network is thus the student model that mimics the motion attention map. With probabilistic attention modeling, the imitation of the attention maps is enforced by using the loss

$$\mathcal{L}^A = \sum_t KL \left[\mathcal{A}^M(t) \parallel \tilde{\mathcal{A}}^M(t) \right]. \quad (6)$$

This loss minimizes the distance between the attention maps at every time step t . In our implementation, our teacher flow network is trained with the same attention mechanism. Once trained, the weights of the teacher model remain fixed during the learning of the student model. At testing time, only the student model (RGB network) is used for inference.

3.5 Our Full Model

Putting everything together, we summarize our full model with probabilistic soft attention and attention distillation. Specifically, our model estimates the two probabilistic attention maps $\mathcal{A}^M \sim F_{\mathcal{A}}^M(\phi(x))$ (motion) and $\mathcal{A}^A \sim F_{\mathcal{A}}^A(\phi(x))$ (appearance). These maps are further used to predict the action labels. This is given by

$$\tilde{y} = F_{\mathcal{R}}^M(\phi(x), \mathcal{A}^M) + F_{\mathcal{R}}^A(\phi(x), \mathcal{A}^A) \quad (7)$$

where each $F_{\mathcal{R}}$ follows Eq 4. We use equal weighting for $F_{\mathcal{R}}^M$ and $F_{\mathcal{R}}^A$. We found that tuning the weights has negligible effect on the performance in practice.

Loss Function. Our training loss is defined as

$$\mathcal{L} = CE(\tilde{y}, y) + \lambda_1 \sum_t KL \left[\mathcal{A}^M(t) \parallel \tilde{\mathcal{A}}^M(t) \right] + \lambda_2 \sum_t KL \left[\mathcal{A}^A(t) \parallel U \right], \quad (8)$$

where CE is the cross entropy loss between the predicted labels \tilde{y} and the ground-truth y . Thus, the loss consists of three terms. The first cross entropy term is to minimize the error for classification. The second KL term (from Eq. 6) enforces that the motion attention \mathcal{A}^M should mimic the attention map $\tilde{\mathcal{A}}^M$ from the reference flow network. Finally, the third KL term (from Eq. 3) regularizes the learning of the appearance attention. The coefficients λ_1 and λ_2 are used to balance the three terms. We choose $\lambda_1 = 1$ and $\lambda_2 = 1/(T_\phi \times W_\phi \times H_\phi)$.

3.6 Implementation Details

Network Architecture. Our model uses I3D network [9] as the backbone. I3D has five 3D convolution blocks, and three of them are composed of multiple Inception Modules. For all attention modules, the intermediate feature ϕ is obtained from the outputs of the 4th convolutional block. The attention map is used to select the final network feature from the last Inception module of the 5th convolutional block.

Data Preparation. We down-sample all frames to 320×256 with a frame rate of 24 Hz. For training, we compute optical flow using TV-L1 [52]. We apply several data augmentation techniques, including random flipping, cropping and color perturbation to prevent overfitting. Our model takes 24 consecutive frames as inputs, and all input frames are cropped to

224×224 for training. For testing, we evaluate our model on full resolution clips (320×256) and aggregate scores from all clips to produce the video-level results.

Training and Inference Details. All of our models are trained using SGD with momentum of 0.9. The weights are initialized from Kinetics pre-trained models provided by the authors of [4]. Our models are trained with a batch size of 64 on 4 GPUs. The initial learning rate is 0.01 with a decay rate of 10 when the loss starts to saturate. We set weight decay to $4e-5$ and enable batch norm [14]. We also adopt dropout rate 0.7. *At inference time our model does not need optical flow*, and runs at the same speed as the RGB network.

4 Experiments

We now present our experiments and results. We start with a systematical evaluation of attention guided action recognition, followed by our main results on several public datasets.

4.1 Datasets and Metrics

We make use of four action recognition datasets for our experiments: UCF101, HMDB51, EGTEA Gaze+ (egocentric videos) and 20BN-V2. UCF101 [68] has 13,320 videos from 101 action categories. HMDB51 [21] includes 6,766 videos from 51 action categories. EGTEA [22, 23] contains 10,321 videos from 106 action categories. We evaluate mean class accuracy and report the results using the first split of these three datasets. 20BN-V2 [29] has over 220K videos from 174 fine-grained action categories. We use their training and validation split, and report top-1/top-5 accuracy.

4.2 Attention Guided Action Recognition

We start from an ablation study of attention-guided action recognition. Specifically, we evaluate different combinations of attention modules and compare their results to those from models without attention. Our experiments show that the proper design of the attention mechanism can consistently improve the performance of action recognition across multiple datasets. We now present our baselines and results.

Baselines. We consider the different combinations of how the model generates attention maps (Soft vs. Probabilistic Attention) and how the attention maps are used for recognition (Attention Pooling vs. Residual Connection). In addition, we also show how the approach to combining motion attention and appearance attention affects the recognition performance. The valid combinations include the following:

- **Soft-Atten** combines soft attention and attention pooling for recognition similar to [26].
- **Soft-Res** is the residual attention in [44] that adds residual connection to Soft-Atten.
- **Prob-Atten** combines probabilistic attention with attention pooling as in [22].

We note that the combination of Prob+Res is invalid, as it violates the probabilistic modeling of attention. In practice, we also found its training to be unstable. Therefore, we report the results of three valid designs for both RGB and flow stream and the vanilla I3D models (our backbone) using the same input sequence length (24 frames) in Table 1. Adding attention to the backbone recognition network almost always improves the performance. Importantly, Soft-Res decreases the performance of RGB stream on HMDB51 and Soft-Atten decreases the performance of RGB stream on HMDB51 and UCF101. More interestingly, Prob-Atten

Method	UCF101	HMDB51	EGTEA
Flow I3D	94.0	73.9	38.3
Flow Soft-Atten	94.7	74.1	39.1
Flow Soft-Res	95.2	74.4	39.5
Flow Prob-Atten	94.9	74.2	40.4
RGB I3D	94.8	70.9	47.3
RGB Soft-Atten	94.7	70.8	48.6
RGB Soft-Res	94.9	70.1	48.6
RGB Prob-Atten	95.1	71.3	49.1

Table 1: Evaluations of attention modules. We compared 3 different design choices with RGB/flow stream on three datasets. Prob-Atten provides a consistent performance boost on both streams and across datasets.

Method	UCF101	HMDB51
Dynamic Image [10]	90.6	61.3
ActionFlowNet [15]	83.9	56.4
TVNet [11]	94.5	71.0
I3D RGB* [9]	94.8	70.9
FeatMatch [18]	94.3	70.7
MARS [8]	94.6	72.3
Ours (Prob-Distill)	95.7	72.0
Two Stream ResNeXt [6]	95.6	74.0
MARS+Flow ResNeXt [8]	94.9	74.5
Two Stream I3D*	96.7	74.8
Prob-Distill+Flow I3D*	97.4	75.7

Table 2: Action recognition results on UCF101 and HMDB51 datasets. We compare the results of our model with previous works. Our model outperforms state-of-the-art methods that use only RGB stream and the same input sequence length by $\sim 1\%$. *For fair comparison, we report results of I3D models that use 24 frames as inputs—the same as our model.

is the most robust design choice, despite the lack of human gaze as a supervisory signal as in [22]. Across all of the modalities and datasets, Prob-Atten can consistently improve the recognition accuracy (+0.3%/0.4%/1.8%) for the RGB stream and (+0.9%/0.5%/2.1%) for the flow stream. The performance boost from the attention module is larger for the flow stream in comparison to the RGB stream. Moreover, attention modules provide more significant boost for egocentric actions (EGTEA). We conjecture that the explicit modeling of attention helps to suppress background objects in first person video.

4.3 Attention Distillation for Action Recognition

We now evaluate our method of attention distillation. In this setting, we assume a reference flow network with attention module is given at training time. We attach motion and appearance attention modules to our RGB backbone. Both attention heads follow the same attention module design as the reference network. The flow attention is asked to mimic the motion attention map from the reference flow network. *During testing our model does not need optical flow*, and runs at the same speed as the RGB network (about 100 times faster than a two stream network [6]). We present our results for action recognition, and contrast our method with feature distillation methods [6, 18].

Impact of Attention Distillation. Table 2 compares our results with previous methods on UCF101/HMDB51. We denote our models using Prob-Atten for distillation as *Prob-Distill*. Prob-Distill outperforms all previous state-of-the-art methods of motion represen-

Method	Top-1/Top-5 Acc	Temporal Footprints
TRN RGB [50]	48.8 / 77.6	5 sec
TRN RGB+Flow [50]	55.5 / 83.1	5 sec
I3D RGB	47.3 / 76.1	1 sec
I3D Flow	46.7 / 75.9	1 sec
Ours (Prob-Distill)	49.9 / 79.1	1 sec
Two Stream I3D	53.7 / 82.5	1 sec
Prob-Distill+Flow I3D	54.6 / 83.0	1 sec

Table 3: Action recognition results on on 20BN-V2 dataset [29]. Our model achieves the best performance among networks that uses RGB frames. Fusing our model with a flow network also outperforms two stream baseline by a significant margin.

tation learning. Specifically, our results are at least 1.2% better than previous state-of-the-art methods for learning motion-aware video representations from RGB frames, including Dynamic Image [11], ActionFlowNet [61] and TVNet [2]. Our model also outperforms MARS [5], our direct competitor, by 0.9% on UCF101 and performs on-par with MARS on HMDB51 when using a similar sequence length, despite the fact that MARS uses a stronger backbone network. It is worth noting that this performance boost is significant for action recognition. In contrast, with 50 more layers, ResNet101 is only 0.7% better than ResNet50 on HMDB51 [15]. Moreover, Prob-Distill also outperforms another feature distillation method – FeatMatch [48] by a significant margin (+1.4%/1.3% on UCF101/HMDB51). These results support our argument that *distilling attention maps is more robust than distilling network features for motion representation learning*. Finally, a late fusion of our model with a reference flow network helps to further boost the performance.

Table 3 presents our results on a large scale dataset—20BN-V2. With 1/5 of the temporal receptive field as TRN [50], our model with RGB frames outperforms TRN RGB by 1.1%/1.5% in top-1/top-5 accuracy. And our method improves the backbone by 2.6%/3.0% in top-1/top-5 accuracy. Further fusion of our model with a flow network improves the results by a large margin (+4.7%), again outperforming the two stream baseline.

Learning from a Weak Flow Network. Crasto et al. [5] pointed out that their model ran into a failure mode when the reference flow network has worse performance than the RGB network. To support our claim that attention distillation can leverage a flow-based teacher network even when the flow network does not provide strong baseline performance, we report the results of our model on the EGTEA Gaze+ dataset. Due to severe ego-motion, flow-based models are less effective than RGB models on this dataset. For instance, I3D Flow is 9% worse than I3D RGB (38.3% vs. 47.3%). Despite a much weaker teacher model, Prob-Distill achieves 49.5%, outperforming the best attention-based I3D models for both RGB (Prob-Atten 49.1%) and Flow (Prob-Atten 40.4%). This indicates that even with a weak teacher model, our proposed method is a robust approach to video representation learning.

Distillation without Forgetting. Feature distillation might “overwrite” the features from RGB stream with the features from flow stream. This is evidenced by the result that fusing MARS with reference flow stream network lags behind the two stream version of the network (MARS + Flow ResNeXt vs. Two Stream ResNeXt in Table 2). In contrast, fusing our Prob-distill model with a reference flow model (Prob-Distill + Flow I3D in Table 2) further improves the accuracy and outperforms the two-stream I3D model (+0.5% on UCF101, +0.7% on HMDB51 and +0.9% on 20BN-V2). These results indicate that our attention distillation model does not simply copy the feature from the reference flow network, as the distilled RGB model can still preserve meaningful appearance features.

Note that our single-stream (Prob-Distill) results still lag behind the two stream networks

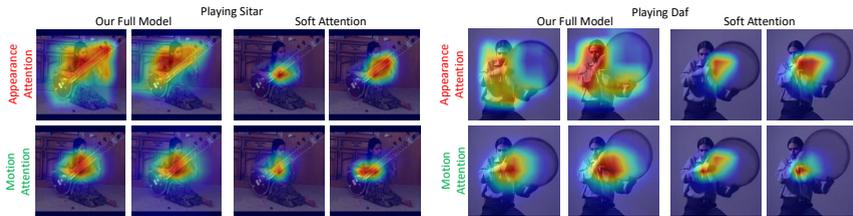


Figure 2: Visualization of attention maps (Ours vs. Soft-Atten using the same I3D backbone). For each video clip, we re-interpolate the attention maps and plot them on the first and last frame. Red regions indicate higher value of attention. Our model produces appearance and motion attention maps that are qualitatively different and index key action regions.

when using the same input sequence length (Two Stream I3D*). This gap reveals that our model does not fully capture the concepts of motion that are encoded in the two stream networks. Nonetheless, we believe that our model provides a key step forward for learning motion-aware representations from RGB frames. Note that some most recent works achieved better performance on the benchmark datasets using more advanced network structure [15, 33, 45, 46], additional features [9], or a longer temporal footprint [3, 5]. In this context, our work provides a novel method for learning video representations and a robust strategy for knowledge distillation. In supplementary materials, we provide additional analysis to show how the learned attention maps help to localize the spatial extent of actions. We also demonstrate how motion information is encoded in the distilled model.

Visualization of Attention Maps. To better understand our model, we visualize both motion and appearance attention maps from our model. We also compare these maps with attention maps created by our Soft-Atten models from RGB and flow streams in Fig 2. Notice that these two attention maps are qualitatively different across all methods. The appearance attention is likely to cover foreground objects or actors, while the motion attention focuses on the moving parts. Moreover, the appearance attention from our model can better localize the foreground regions of actions than those of Soft-Atten from the RGB stream, while the motion attention from our model remains similar to the Soft-Atten from the flow stream. We also find that the attention maps from our model are more “diffused”. This is because the regularization by a uniform distribution in Prob-Atten leads to smoother attention maps.

5 Conclusions

In this paper, we presented a novel method of attention distillation for action recognition. We provided extensive experiments to evaluate our method. Our results demonstrate that proper design of the attention module helps to improve recognition performance. In addition, attention maps from RGB and flow networks are qualitatively different, suggesting that these networks capture different aspects of the video. We also showed that our method achieves competitive results for action recognition across datasets, and that attention distillation is more robust for learning a motion-aware video representation. We believe our work provides valuable insights into attention based recognition, and a step towards learning spatio-temporal video features via knowledge distillation.

Acknowledgments. Portions of this research were supported in part by National Science Foundation Award 1936970 and a gift from Facebook. YL acknowledges the support from the Wisconsin Alumni Research Foundation. XC acknowledges the support from Midea Emerging Technology Co., Ltd.

References

- [1] Hakan Bilen, Basura Fernando, Efstratios Gavves, and Andrea Vedaldi. Action recognition with dynamic image networks. *TPAMI*, 2018.
- [2] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *SIGKDD*, 2006.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [4] Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, and Cordelia Schmid. Potion: Pose motion representation for action recognition. In *CVPR*, 2018.
- [5] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. MARS: Motion-Augmented RGB Stream for Action Recognition. In *CVPR*, 2019.
- [6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *ECCV*, 2018.
- [7] Lijie Fan, Wenbing Huang, Stefano Ermon Chuang Gan, Boqing Gong, and Junzhou Huang. End-to-end learning of motion representation for video understanding. In *CVPR*, 2018.
- [8] Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes, and Andrew Zisserman. What have we learned from deep representations for action recognition? In *CVPR*, 2018.
- [9] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 1999.
- [10] Nuno Garcia, Pietro Morerio, and Vittorio Murino. Modality distillation with multiple stream networks for action recognition. In *ECCV*, 2018.
- [11] Rohit Girdhar and Deva Ramanan. Attentional pooling for action recognition. In *NeurIPS*, 2017.
- [12] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *CVPR*, 2019.
- [13] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *CVPR*, 2016.
- [14] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*, 2018.
- [15] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet. In *CVPR*, 2018.
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

- [18] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.
- [19] Boyuan Jiang, MengMeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. Stm: Spatiotemporal and motion encoding for action recognition. In *ICCV*, 2019.
- [20] 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.
- [21] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: a large video database for human motion recognition. In *ICCV*, 2011.
- [22] Yin Li, Miao Liu, and James M. Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *ECCV*, 2018.
- [23] Yin Li, Miao Liu, and James M Rehg. In the eye of the beholder: Gaze and actions in first person video. *arXiv preprint arXiv:2006.00626*, 2020.
- [24] Zhenyang Li, Kirill Gavrilyuk, Efstratios Gavves, Mihir Jain, and Cees GM Snoek. Videolstm convolves, attends and flows for action recognition. *CVIU*, 2018.
- [25] Miao Liu, Siyu Tang, Yin Li, and James Rehg. Forecasting human object interaction: Joint prediction of motor attention and actions in first person video. In *ECCV*, 2020.
- [26] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *CVPR*, 2019.
- [27] Zelun Luo, Jun-Ting Hsieh, Lu Jiang, Juan Carlos Niebles, and Li Fei-Fei. Graph distillation for action detection with privileged modalities. In *ECCV*, 2018.
- [28] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017.
- [29] Farzaneh Mahdisoltani, Guillaume Berger, Waseem Gharbieh, David Fleet, and Roland Memisevic. Fine-grained video classification and captioning. *arXiv preprint arXiv:1804.09235*, 2018.
- [30] Volodymyr Mnih, Nicolas Heess, Alex Graves, and Koray Kavukcuoglu. Recurrent models of visual attention. In *NeurIPS*, 2014.
- [31] Joe Yue-Hei Ng, Jonghyun Choi, Jan Neumann, and Larry S Davis. Actionflownet: Learning motion representation for action recognition. In *WACV*, 2018.
- [32] Javier Sánchez Pérez, Enric Meinhardt-Llopis, and Gabriele Facciolo. TV-L1 optical flow estimation. *IPOL*, 2013.
- [33] AJ Piergiovanni and Michael S. Ryoo. Representation flow for action recognition. In *CVPR*, 2019.
- [34] Ronald Poppe. A survey on vision-based human action recognition. *Image and vision computing*, 2010.

- [35] Shikhar Sharma, Ryan Kiros, and Ruslan Salakhutdinov. Action recognition using visual attention. In *ICLR Workshop*, 2016.
- [36] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016.
- [37] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014.
- [38] K. Soomro, A. Roshan Zamir, and M Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. In *CRCV-TR-12-01*, 2012.
- [39] Jonathan Stroud, David Ross, Chen Sun, Jia Deng, and Rahul Sukthankar. D3d: Distilled 3d networks for video action recognition. In *WACV*, 2020.
- [40] Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, and Wei Zhang. Optical flow guided feature: A fast and robust motion representation for video action recognition. In *CVPR*, 2018.
- [41] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [42] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [44] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, 2017.
- [45] Jue Wang, Anoop Cherian, Fatih Porikli, and Stephen Gould. Video representation learning using discriminative pooling. In *CVPR*, 2018.
- [46] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018.
- [47] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [48] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017.
- [49] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, 2019.
- [50] Bolei Zhou, Alex Andonian, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018.