

MOD: A Deep Mixture Model with Online Knowledge Distillation for Large Scale Video Temporal Concept Localization

Rongcheng Lin, Jing Xiao, Jianping Fan
University of North Carolina at Charlotte

{rlin4, xiao, jfan}@uncc.edu

Abstract

In this paper, we present and discuss a deep mixture model with online knowledge distillation (MOD) for large-scale video temporal concept localization, which is ranked 3rd in the 3rd YouTube-8M Video Understanding Challenge. Specifically, we find that by enabling knowledge sharing with online distillation, finetuning a mixture model on a smaller dataset can achieve better evaluation performance. Based on this observation, in our final solution, we trained and finetuned 12 NeXtVLAD models in parallel with a 2-layer online distillation structure. The experimental results show that the proposed distillation structure can effectively avoid overfitting and shows superior generalization performance. The code is publicly available at: https://github.com/linrongc/solution_youtube8m_v3

1. Introduction

Temporal concept localization within videos, which aims at automatically recognizing/retrieving topic related video segments, is one of critical and challenging problems to enable real world applications, including video search, video summarization, action recognition and video content safety etc. To accelerate the pace of research in this area, Google Research launched the 3rd YouTube-8M video understanding challenge and released 237K human-verified segment labels in addition to the about 6M noisy video-level labels. The goal is to retrieve related video segments from an unlabeled testing set for each of the 1000 classes. How to effectively leverage the large but noisy video-level labels for temporal localization is the main challenge.

One of the straightforward ideas is to pretrain models on the video-level dataset and then finetune the models using the smaller segment-level dataset. This approach turns out to be very effective in solving the problem. Also, we find that increasing parameter number of models by making the model wider can further improve the performance. But the

marginal gains quickly diminish as the model are more likely to overfit the training dataset. Another way to increase the complexity of prediction system is to combine multiple models. Techniques to combine a set of weaker learners to create a strong learner, including bagging and boosting, are widely used in solving traditional machine learning problems. It is capable of reducing model variance and avoiding overfitting. However, in the era of deep learning, with millions even billions of parameters, single neural network could easily overfit the whole training dataset. The marginal gains from naive ensemble of multiple similar models also quickly diminish.

In this work, we propose a new approach by training a mixture of multiple base models in parallel with online knowledge distillation. With similar parameter number, a mixture model with online knowledge distillation can generalize better in the finetuning task than the wider model or the naive mixture. One possible explanation is that the online distillation part give each of the base models a holistic view of the similarity space and avoid the mixture model to overfit the smaller dataset. Based on this assumption, we built a 2-layer mixture model, which is a mixture of 4 MixNeXtVLAD models. And each of the MixNeXtVLAD model is a mixture of 3 base NeXtVLAD models[16]. In summary, we trained 12 NeXtVLAD models in parallel and enabled a 2-layer online distillation structure. Experimental results show the superior generalization performance in the finetuning task, compared to simple mixture models or wider models.

2. Related Work

2.1. Deep Neural Network for Video Classification

With the availability of large-scale video dataset, researchers proposed many deep neural networks and achieved remarkable advances in the field of video classification. In general, these approaches can be roughly summarized into 4 categories: (a) **Spatiotemporal network**[13][12][25]. By regarding the temporal dimension as the extension of spatial dimensions, these models main-

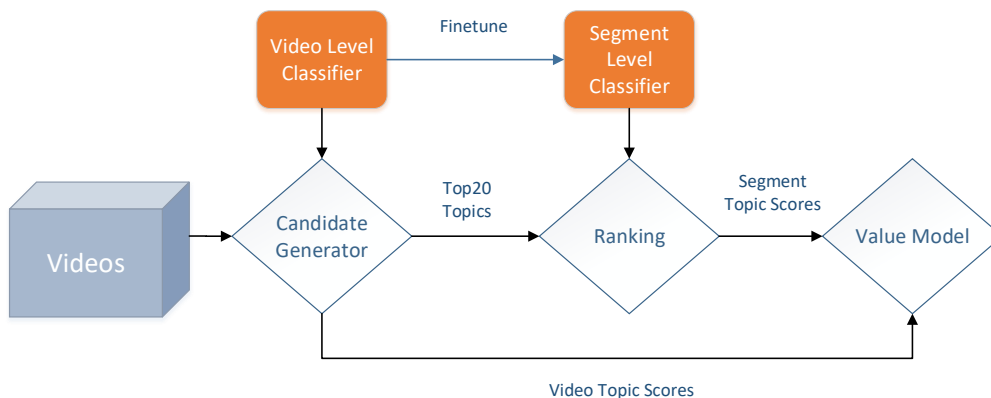


Figure 1. Solution overview.

ly rely on 2D or 3D convolution and pooling to aggregate information in the videos. (b) **Recurrent network**[4][5]. Apply recurrent neural networks, such as LSTM and GRU to aggregate the sequential actions in the videos. (c) **Two Stream Network**[22][8][28][18]. Utilize optical flow images or similar features to model the motion in the video separately. The features extracted from frame images network and the optical flow network are fused to represent the videos. (d) **Other approaches**[9][27][6][26]. Use other information or methods to generate features for video representation and classification.

2.2. Learnable Pooling Methods

In the field of computer vision, aggregating multiple features into a single compact feature vector has been a long-standing research problem. Techniques, including BoW(Bag of visual Words)[23], FV(Fisher Vector)[19] and VLAD(Vector of Locally Aggregated Descriptors)[11], are widely used in computer vision systems[14][21], including image/video retrieval, classification and localization. Recently, inspired by the work of VLAD, a learnable pooling network, NetVLAD, is firstly introduced in [3] to solve the problem of place recognition. In the task of video understanding with pre-extracted frame level features, NetVLAD shows superior performance[3][17]. Several other variants, including NeXtVLAD[16] and Non-local NetVLAD[24] etc, were proposed to further improve the parameter efficiency and generalization performance.

2.3. Knowledge Distillation

Knowledge distillation[10] is an effective and popular approach for model compression by distilling a complex teacher model to a simpler student model. The success of transferring the dark knowledge between networks has inspired many novel research work in com-

puter vision[15][20][7]. Recently, researchers find that, rather than the one-way knowledge transfer, enabling collaborative learning of several simple student models with a two-way knowledge sharing can achieve superior results[30][29] and can be efficiently trained within a distributed training system[2].

3. Solution

3.1. Solution Overview

The overall structure of the solution to generate video segments for each of the 1000 topics is illustrated in Figure 1. The system is comprised of three phases: (1) candidate generation via a video level classifier. Only the top 20 topics are considered to be existed in the video. An offline analysis demonstrate that those candidates cover over 97% of the positive samples(recall) in the segment training dataset. This step significantly reduce the search space. (2) a segment level classifier is used as a ranker to assign probabilities to each of the 5s segments in the video. The segment level classifier is directly finetuned from the video level classifier. (3)we combine the video topic scores $P_{vid}(K)$ and segment topic score $P_{seg}(K)$ via a value model:

$$P(K = k) = P_{vid}^{0.05}(K = k) * P_{seg}^{0.95}(K = k) \quad (1)$$

And finally, for each of the 1000 topics, we retrieve the top 10K video segments ranked by the combined score. The whole system relied heavily on the performance of video and segment level classifier. How to build a accurate and robust classifier is the essential part of the solution.

3.2. NeXtVLAD Model

The base model used in our classifier is NeXtVLAD model, which achieved the best single model performance




Figure 2. NeXtVLAD pooling network.

in the 2nd YouTue-8M video understanding challenge. Before diving into the final solution, we will first briefly review the NeXtVLAD pooling network and NeXtVLAD model for video classification.

A NeXtVLAD pooling network, as shown in Figure 2, is a variant of NetVLAD, which is a differentiable network inspired by traditional Vector of Locally Aggregated Descriptors(VLAD). Considering a video input x with M frames and each of the frame is represented as N -dimension feature, a NeXtVLAD expand the input dimension by a factor of λ at first via a linear projection to be \hat{x} with a shape of $(M, \lambda N)$. Then \hat{x} is splitted into G groups, each of which is represented as \hat{x}^g . The NeXtVLAD pooling is a mixture of group-level NetVLAD aggregations into K clusters:

$$y_{jk} = \sum_g \alpha_g(\hat{x}_i) v_{jk}^g \quad (2)$$

$$\forall i \in \{1, \dots, M\}, j \in \{1, \dots, N\}, k \in \{1, \dots, K\}, g \in \{1, \dots, G\}$$

in which $\alpha_g(\hat{x}_i)$ is group level attention function:

$$\alpha_g(\hat{x}) = \sigma(w_g^T \hat{x} + b_g) \quad (3)$$

and v_{jk}^g is the output of group level NetVLAD aggregation:

$$v_{jk}^g = \sum_i \alpha_{gk}(\hat{x}_i) (\tilde{x}_{ij}^g - c_{kj}) \quad (4)$$

Finally, a l_2 normalization, a.k.a. intra-normalization, is applied to the aggregated features for each of the clusters:

$$\hat{y}_{jk} = \frac{y_{jk}}{\|y_k\|_2} \quad (5)$$

The l_2 normalization is one of essential parts to make features extracted from different videos or video segments are comparable. And it is also one of the reasons why finetuning a video-level model can work well as a segment level classifier. As illustrated in Figure 5, in the NeXtVLAD model designed for video classification, video and audio features are aggregated by two NeXtVLAD pooling networks separately. Then the aggregated features are concatenated and fed into a dropout layer before a FC layer

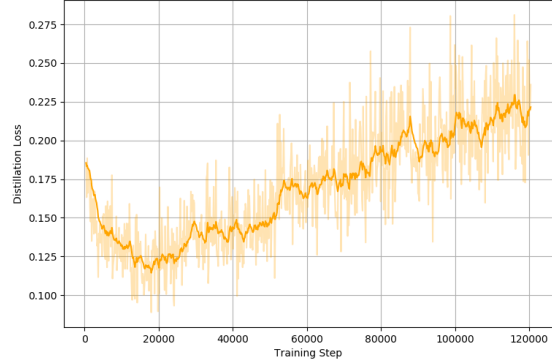


Figure 3. An example of distillation loss of a MixNeXtVLAD Model.




Figure 4. An example of label loss and regularization loss of a MixNeXtVLAD Model.

is applied to reduce the dimension of the encoded features. After the dropout layer, a context gating layer is appended to capture the dependency among topics. Finally, a logistic model is used as the final classifier.

3.3. MixNeXtVLAD Model

Training multiple base models in parallel and distill knowledge from the mixture predictions to sub-models via a distillation loss is firstly introduced in [29] and applied to the video classification problem in [16]. The MixNeXtVLAD model is a mixture of 3 NeXtVLAD model with on-the-fly knowledge distillation. As shown in Figure 6, the logit z^e of mixture prediction p^e is the weighted sum of the logits z^m from predictions p^m of sub-models. Given the ground truth label y , The final loss of the MixNeXtVLAD model is:

$$\mathcal{L} = \sum_{m=1}^3 \mathcal{L}_{bce}(y, p_m) + \mathcal{L}_{bce}(y, p_e) + T^2 * \sum_{m=1}^3 \mathcal{KL}(Soft(p_e, T) \| Soft(p_m, T)) \quad (6)$$

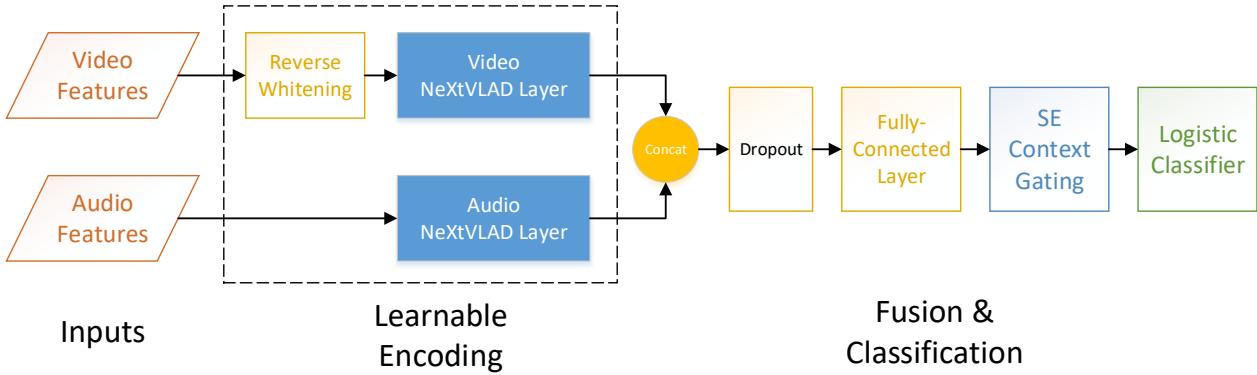


Figure 5. Overview of a NeXtVLAD model designed for Youtube-8M video classification.

where L_{bce} is the binary cross entropy and $\mathcal{KL}(Soft(p_e, T) || Soft(p_m, T))$ represents distillation loss, which is the KL divergence between the softened predictions:

$$Soft(p, T) = Softmax(z/T) \quad (7)$$

in which z is the logits of prediction p . A larger T value will emphasize more on the smaller values in the prediction and thus share more knowledge about the learned similarity in the task space.

One of the main assumptions is that the online distillation loss will provide a holistic view for sub-models to the task space during training. If we dive closer to the binary cross entropy loss of the mixture prediction $\mathcal{L}_{bce}(y, p_e)$, we can find the loss capture the remaining part which is not covered by the predictions from all the sub-models. In other word, if one sub-model capture part of the true prediction, then the information will be ignored by the loss for other sub-models. As a result, the predictions of sub-models are diversified during training. The distillation loss between mixture prediction and individual prediction will ensure the sub-models have the holistic view of the whole task space. Figure 3 shows one example of the online distillation loss of a MixNeXtVLAD Model during training. The distillation loss is optimized(decreasing) at the beginning then increasing steadily as we further minimize the whole objective function (Figure 4). The increase of the distillation loss at later stage of training is a implicit proof of our assumption.

3.4. Deep Mixture of NeXtVLAD Models with Online Distillation

A deep mixture of NeXtVLAD models with online distillation (MODNeXtVLAD thereafter), which is the model used as our final solution, is a intuitive extension of the MixNeXtVLAD Model. As shown in Figure 7, MOD-

NeXtVLAD is a mixture of 4 MixNeXtVLAD models, each of which is a mixture of 3 base NeXtVLAD Models. So in total, in MODNeXtVLAD, 12 NeXtVLAD models are trained and finetuned simultaneously. As for the knowledge distillation part, knowledge is firstly distilled from the final prediction to each of the mixture models, then from mixture prediction to each of the NeXtVLAD models. For simplicity, we apply the same parameter(T in this case) in the two-stage knowledge distillation.

To be general, the MOD structure forms a simple 2-layer model-level hierarchy, where each sub-tree is an independent mixture model and knowledge is distilled from root to leaves one layer at a time. One advantage of the MOD structure is its suitability for distributed training. Except for knowledge distillation loss and mixture of logits, models in different subtrees can be trained independently and thus can be located in different physical devices and the communication(network) overhead is negligible. So in our implementation, we applied model parallel distributed training strategy instead of data parallel to improve the training efficiency. With data parallel strategy, the training speed of one NeXtVLAD model in 2 Nvidia 1080TI GPUs is about 400 examples per second. By enabling model parallel in training 12 same NeXtVLAD models with MOD structure and model parallel strategy, we can achieve a training speed of 140+ examples per second using 4 Nvidia 1080 TI GPUs.

4. Experimental Results

4.1. Dataset and Evaluation Metrics

4.1.1 YouTuebe-8M Video Dataset

Youtube-8M video dataset[1] consists of 6.1M popular videos from Youtube.com. These videos are splitted in to 3 partitions: training(70%), validation(20%) and testing(10%). For each video in the training and validation dataset, one or multiple labels(3.0 labels/video on average)

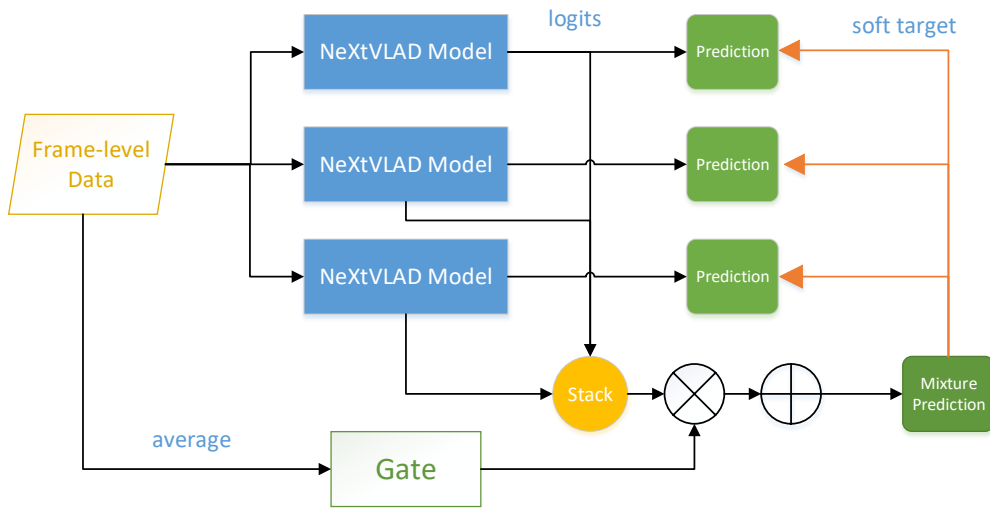


Figure 6. Overview of a mixture of 3 NeXtVLAD models(MixNeXtVLAD) with online knowledge distillation. The orange arrows indicate the distillation of knowledge from the mixture prediction to the predictions of sub-models.

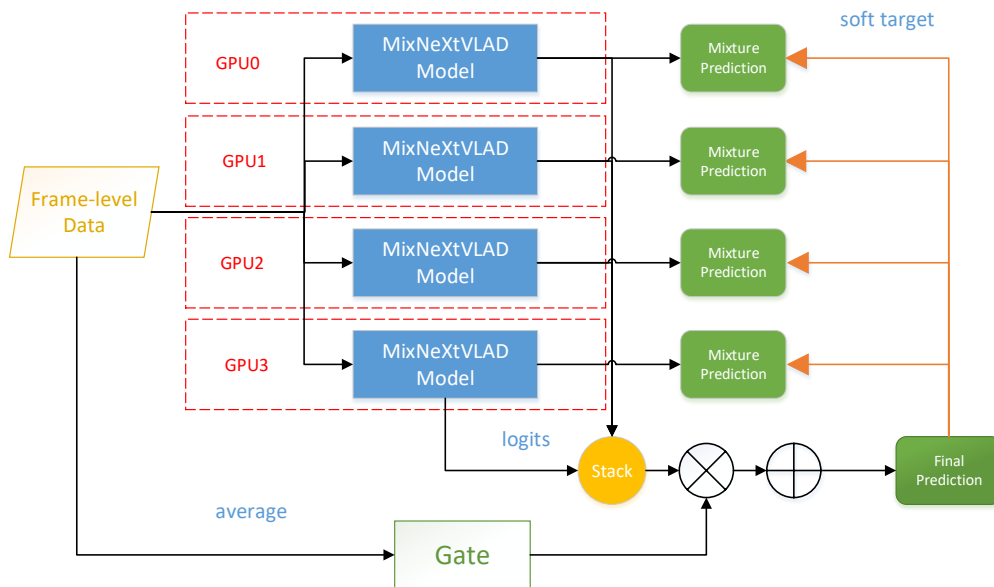


Figure 7. Overview of a mixture of 4 MixNeXtVLAD models with online knowledge distillation. Each of the MixNeXtVLAD model is a mixture of 3 base NeXtVLAD model. The orange arrows indicate the distillation of knowledge from the final prediction to mixture predictions of subordinated mixture models. Each of the 4 MixNeXtVLAD model is located in the separate GPU.

are generated by an annotation machine from a vocabulary of 3862 visual entities. These video-level labels is not verified by human and thus noisy in terms of label quality. For every second of the videos, frame-level features, including a 1024-dimensional visual feature and a 128-dimensional audio feature, are precomputed and provided for model training.

4.1.2 Youtube-8M Segment Dataset

As an extension of the original YouTube-8M dataset, the segment dataset contains 237K human-verified segment labels on 1000 classes. These segments are sampled from the validation set of the Youtube-8M video dataset and contains exactly 5 frames. Each segment label indicates whether the 5s segments contains objects of the target class. Compared

Table 1. Performance comparison of single models trained on frame-level features. The parameters inside the parenthesis represents (group number G, cluster number K, expansion factor λ , hidden size H).

	Model	Parameter	Private LB	Public LB
Without Finetune	Dummy prediction	83M	0.64809	0.66188
Without Pretrain	DBoF Baseline	16M	0.69882	0.71077
	NeXtVLAD(8G, 128K, X2, 2048H)	83M	0.77009	0.77730
Pretrain and finetune	NeXtVLAD(8G, 128K, X2, 2048H)	83M	0.79642	0.80635
	NeXtVLAD_large(8G, 256K, X4, 2048H)	320M	0.80586	0.81611
	NeXtVLAD_distill(8G, 128K, X2, 2048H)	83M	0.81509	0.82267

Table 2. Performance comparison of mixture models. All the base models used are NeXtVLAD(8G, 128K, X2, 2048) except for MixNeXtVLAD_large, which take NeXtVLAD_large as the base model.

	Model	Base Model Number	Parameter	Private LB	Public LB
One-Layer Mixture	MixNeXtVLAD(T=0)	3	250M	0.80797	0.81688
	MixNeXtVLAD(T=1)	3	250M	0.81125	0.82023
	MixNeXtVLAD(T=10)	3	250M	0.81617	0.82477
	MixNeXtVLAD(T=20)	3	250M	0.81984	0.82699
	MixNeXtVLAD_large(T=20)	4	1280M	0.82262	0.83014
Two-Layer Mixture	MODNeXtVLAD(T=20)	12	1000M	0.82512	0.83251

to the video dataset, this segment dataset is clean but much smaller. How to leverage the large amount but noisy video level labels is one of the main challenges.

4.1.3 Evaluation Metrics

In the 3rd Youtube-8M video understanding challenge, submissions are evaluated using Mean Average Precision at K (MAP@K):

$$MAP@K = \frac{1}{C} \sum_{c=1}^C \frac{\sum_{k=1}^K P(k) * rel(k)}{N_c} \quad (8)$$

where C is the number of classes, N_c is the total positive samples in the class, $P(k)$ is the precision at cutoff k and $rel(k)$ is an indicator function to represent whether the k_{th} items belong to class c . The metric is an approximate of the area under Precision-Recall curve.

4.2. Implementation Details

Our implementation is based on the TensorFlow starter code provided by the organizer. All the models are run at a machine with 4 Nvidia GPUs. We follows the same settings in [] to train video level models. For a fair comparison, each model is trained for about 500K steps to guarantee the convergence. As for larger models, including our final model, we use a batch size of 80 to avoid out of memory in GPUs.

In the finetuning stage, all the models are trained with a batch size of 512. The dropout rate and the l2-normalization penalty are increased to 0.75 and 1e-4 respectively aiming to prevent overfitting. Models are trained for 10 epochs on the segment dataset using the Adam optimizer with a initial learning rate of 0.0002. The learning rate is decayed by a

factor of 0.8 for every 1M examples. More training details can be found at https://github.com/linrongc/solution_youtube8m_v3

4.3. Model Evaluation

4.3.1 Single Model Comparison

The performance and parameter number of single models are summarized in Table 1. The evaluation metrics presented in the table is MAP@100000. The models included in the comparison are:

- Dummy prediction. A NeXtVLAD model trained only using the video level labels. All the segments in the video are considered to contains the same content.
- DBoF Baseline. A deep bag of frame model provided in the starter code with 2048 clusters and a hidden size of 1024. The final classifier is a MOE(mixture of experts) model with 5 experts.
- NeXtVLAD. The best single model in the 2nd YouTube-8M video understanding challenge.
- NeXtVLAD_distill. One single NeXtVLAD model used in the two-layer mixture model with online knowledge distillation. It is trained with other 11 NeXtVLAD models with the same settings. But in the inference stage, those 11 NeXtVLAD models are removed.

Generally speaking, models which are pretrained on the larger video dataset outperform models without pretrain.

While a larger NeXtVLAD model with more parameters can achieve better MAP score, one single and small NeXtVLAD model used in 2-layer mixture with online knowledge distillation shows superior performance.

4.3.2 Mixture Model Comparison

We evaluate the one-layer MixNeXtVLAD model (Figure 6) with different settings and one two-layer mixture of NeXtVLAD model, which is the model used in our final submission. As illustrated in Table 2, the MixNeXtVLAD model without knowledge distillation ($T=0$) shows the similar performance with the larger NeXtVLAD model (NeXtVLAD_large). By gradually increasing the value of T , the generalization performance is improved accordingly. The results indicate that, with higher value of temperature (T), more knowledge are distilled from the mixture model to each single model. The knowledge distillation part can effectively avoid model overfitting.

Also, a two-layer mixture model, MODNeXtVLAD, can easily outperform one-layer mixture model even with less number of parameters. The results directly prove the parameter efficiency and better generalization performance of the proposed deep mixture structure with online distillation.

5. Conclusion and Future Work

In this paper, we proposed a novel deep mixture model with online knowledge distillation and evaluated the model performance in the 3rd YouTube-8M video understanding challenge. The model can be efficiently trained in a distributed training system because of the low communication cost between the base models. The experimental results show that, in a finetune task, online knowledge distillation can effectively improve the generalization performance of the mixture model.

Due to the resource limit, only a 2-layer mixture model with online distillation is included in the experiment. Whether a deeper mixture model with online knowledge distillation can further improve the generalization performance still need to be verified.

References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *CoRR*, abs/1609.08675, 2016.
- [2] Rohan Anil, Gabriel Pereyra, Alexandre Passos, Róbert Ormándi, George E. Dahl, and Geoffrey E. Hinton. Large scale distributed neural network training through online distillation. *CoRR*, abs/1804.03235, 2018.
- [3] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [4] Moez Baccouche, Franck Mamalet, Christian Wolf, Christophe Garcia, and Atilla Baskurt. Sequential deep learning for human action recognition. In *Proceedings of the Second International Conference on Human Behavior Understanding*, HBU'11, pages 29–39. Springer-Verlag, 2011.
- [5] Nicolas Ballas, Li Yao, Chris Pal, and Aaron C. Courville. Delving deeper into convolutional networks for learning video representations. *CoRR*, 2015.
- [6] Hakan Bilen, Basura Fernando, Efstratios Gavves, and Andrea Vedaldi. Action recognition with dynamic image networks. *CoRR*, abs/1612.00738, 2016.
- [7] Tianqi Chen, Ian J. Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. In *4th International Conference on Learning Representations, ICLR*, 2016.
- [8] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. *CoRR*, 2016.
- [9] Basura Fernando, Efstratios Gavves, Jose M. Oramas, Amir Ghodrati, and Tinne Tuytelaars. Modeling video evolution for action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [11] Herve Jegou, Matthijs Douze, Cordelia Schmid, and Patrick Prez. Aggregating local descriptors into a compact image representation. In *CVPR*, pages 3304–3311. IEEE Computer Society, 2010.
- [12] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 221–231, Jan. 2013.
- [13] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [14] Ivan Laptev, Marcin Marszaek, Cordelia Schmid, and Benjamin Rozenfeld. Learning realistic human actions from movies. In *IN: CVPR.*, 2008.
- [15] Zhizhong Li and Derek Hoiem. Learning without forgetting. *CoRR*, 2016.
- [16] Rongcheng Lin, Jing Xiao, and Jianping Fan. Nextvlad: An efficient neural network to aggregate frame-level features for large-scale video classification. *CoRR*, abs/1811.05014, 2018.
- [17] David Austin Miha Skalic. Building a size constrained predictive models for video classification. In *ECCV Workshops*, 2018.
- [18] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Computer Vision and Pattern Recognition*, 2015.
- [19] Florent Perronnin and Christopher R. Dance. Fisher kernels on visual vocabularies for image categorization. In *CVPR*. IEEE Computer Society, 2007.

- [20] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *CoRR*, abs/1606.04671, 2016.
- [21] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: A local svm approach. In *Proceedings of the Pattern Recognition, 17th International Conference on (ICPR'04) Volume 3 - Volume 03*, ICPR '04, pages 32–36. IEEE Computer Society, 2004.
- [22] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'14, pages 568–576. MIT Press, 2014.
- [23] J. Sivic and A. Zisserman. Video Google: A text retrieval approach to object matching in videos. In *IEEE International Conference on Computer Vision*, volume 2, pages 1470–1477, 2003.
- [24] Yongyi Tang, Xing Zhang, Jingwen Wang, Shaoxiang Chen, Lin Ma, and Yu-Gang Jiang. Non-local netvlad encoding for video classification. *CoRR*, abs/1810.00207, 2018.
- [25] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, pages 4489–4497. IEEE Computer Society, 2015.
- [26] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. Technical report, arXiv, 2017.
- [27] Xiaolong Wang, Ali Farhadi, and Abhinav Gupta. Actions ~ transformations. In *CVPR*, 2016.
- [28] Zuxuan Wu, Yu-Gang Jiang, Xi Wang, Hao Ye, Xiangyang Xue, and Jun Wang. Fusing multi-stream deep networks for video classification. *CoRR*, 2015.
- [29] Shaogang Gong Xu Lan, Xiatian Zhu. Knowledge distillation by on-the-fly native ensemble. In *arXiv:1806.04606*, 2018.
- [30] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. *CoRR*, 2017.