

Label Denoising with Large Ensembles of Heterogeneous Neural Networks

Pavel Ostyakov, Elizaveta Logacheva, Roman Suvorov, Vladimir Aliev, Gleb Sterkin, Oleg Khomenko, and Sergey I. Nikolenko

Samsung AI Center, Moscow, Russia

Abstract. Despite recent advances in computer vision based on various convolutional architectures, video understanding remains an important challenge. In this work, we present and discuss a top solution for the large-scale video classification (labeling) problem introduced as a Kaggle competition based on the YouTube-8M dataset. We show and compare different approaches to preprocessing, data augmentation, model architectures, and model combination. Our final model is based on a large ensemble of video- and frame-level models but fits into rather limiting hardware constraints. We apply an approach based on knowledge distillation to deal with noisy labels in the original dataset and the recently developed mixup technique to improve the basic models.

Keywords: Video processing, learning from noisy labels, attention-based models, recurrent neural networks, deep learning

1 Introduction

Video understanding and learning high-quality latent representations for videos is an important problem which largely remains open and needs more advances in computer vision, including video processing and scene understanding, audio processing and speech recognition, and natural language processing. One of the first steps towards true video understanding is video classification/labeling, where the problem is to apply labels from a predefined set to a video. While, e.g., text classifiers achieve very high results despite the fact that general language understanding remains a hard problem, for video data even classification is relatively underexplored.

In this work, we concentrate on video classification task presented in the second YouTube-8M Challenge, where the problem is to automatically annotate YouTube videos with a large number of predefined labels based on video-level and frame-level features. While we have introduced several novel modifications to existing deep learning models and performed a large-scale experimental comparison of convolutional, recurrent, and attention-based architectures, we believe that our main contribution lies in dealing with noisy labels. One of the main challenges in the YouTube-8M dataset and other large-scale datasets is that the labels, while relatively high-quality, are not completely reliable. In fact, according to the authors of the dataset, precision and recall of YouTube-8M labels is 78.8%

and 14.5% respectively compared to the gold standard of human raters [1]. In this work, we deal with this problem with the recently developed knowledge distillation approach which we base on a large ensemble of a wide variety of video- and frame-level models. The ensemble, however, is used only for preparing soft labels, and the final model trained on these soft labels is relatively simple and fits into rather limiting hardware requirements imposed in the challenge. Our final resulting model achieved very high results in video classification and placed 4th in the second YouTube-8M Challenge.

The paper is organized as follows. In Section 2, we describe the dataset, introduce the primary evaluation metric that we use throughout the paper, and discuss different data representation methods that we compared for frame-level data. Section 3 describes the main neural architectures used in this work and presents the ensemble-based distillation that we used to alleviate the problem of noisy labels. Section 4 presents an extensive evaluation study of the proposed models and a detailed error analysis, and Section 5 concludes the paper.

2 Data and Evaluation

Dataset. The YouTube-8M dataset contains 5.6M videos encoded as a hidden representation produced by Deep CNN pretrained on the ImageNet dataset [2] for both audio spectrogram and video frames taken at rate of 1Hz (once per second). The dataset also contains aggregated *video-level features* extracted as averaged frame-level features. Each video was automatically annotated with 3862 classes by the YouTube video annotation system. To ensure high quality of the videos in the dataset, videos were preselected to meet the following requirements: each video must be public, have at least 1000 views, be between 120 and 500 seconds long, and also be associated with at least one entity from the target vocabulary; adult and sensitive content has been removed based on automated classifiers. The dataset is an improved version of the initially released YouTube-8M [1], with some noisy and rare labels filtered out, but the labeling still remains rather noisy.

Evaluation. The primary evaluation metric for our experiments and the competition itself is the *global average precision* (GAP), defined as $\text{GAP}@n = \sum_{i=1}^N p(i)r(i)$, where N is the total number of final predictions, $p(i)$ is the precision score at rank i , and $r(i)$ denotes the relevance of prediction i : its 1 if the i -th prediction is correct, and 0 otherwise. For every video top n predicted labels are scored, so $N = nV$, where V is the number of videos.

Frame-level augmentation. Alongside with video-level features, the YouTube-8M dataset contains frame-level features: one preprocessed frame per second for every video. To use frame-level data in our models, we have tested several augmentation strategies. First, we computed and used as features various frame-level statistics, including the mean, standard deviation, median, minimum, maximum,

mode, and video length (number of frames). These statistics were concatenated and used as input for our network architectures.

To use frame-level features in the models, we computed the above statistics for different frames in the video. There are too many frames in a video, and neighboring frames are generally very similar, so it is detrimental to use all frame-level data and preferable to include only a few representative frames. We have tried several subsampling strategies to choose which frames to include as input:

- (1) subsample a subsequence of frames at random or at regular intervals;
- (2) assuming that not all frames are equally important for training, we decided to take one frame per “scene” in the video; for this purpose, we cut the video into “scenes” by thresholding cosine distances between neighbouring frames (Fig. 1 shows sample correlations between the frames in a video; the scene structure is usually rather clear) and take one frame per scene;
- (3) apart from “scenes” specific for each video, we also found global centroids of frame-level features using large-scale k -means clustering for all frames in the dataset, and then used frame statistics from unique nearest centroids for the frames in a given video; with 10,000 centroids, this approach yielded a GAP score of 0.79 (row 14 of Table 1).

Another approach we tried for frame-level information was to try to capture just the *dynamics* in a video sequence. The intuition behind this approach is that the main information about the video may be captured in what is changing between frames and not in the static component. To this end, we have tried several different techniques:

- (1) subtracting per-video average from every frame in the video;
- (2) applying a high-pass filter using the $db3$ wavelet [3];
- (3) computing sparse component with Principal Components Pursuit [4].

However, none of these techniques yielded a GAP score greater than 0.5.

Finally, the best technique that helped us incorporate frame-level information into video-level vectors in the best possible way is as follows. We have taken a number of linear combinations of frames with trainable coefficients (coefficients are linked to frame position) and then applied a ResNet-like model (see Section 3) to the concatenated vectors. This approach yielded a GAP score of 0.855 for 5 linear combinations.

3 Models

General approach. In this section, we define and discuss the models that we have tried in our experiments. In the YouTube-8M competition, one of the requirements for the final model was to fit it into 1Gb of memory. This is a rather limiting constraint for modern state of the art neural networks, so the objective was to both achieve the best possible results in terms of GAP and do

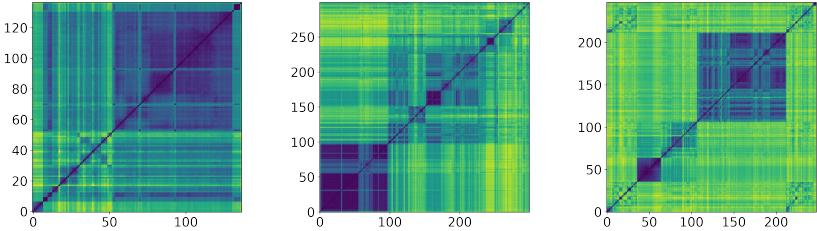


Fig. 1. Cosine distance between subsequent frames in three sample videos; we took abrupt changes in neighboring frames as evidence of a new “scene” beginning.

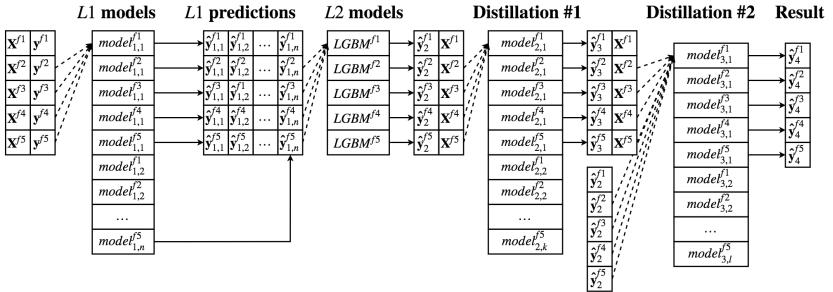


Fig. 2. General flowchart of our approach.

it in limited space (and hence also limited inference time). To achieve both, we propose the following approach, illustrated on Figure 2.

First, we train various models on the video-level and frame-level features provided in the dataset, usually optimizing binary cross-entropy (BCE) with respect to the original hard labels; we call them *first-level* models. The final first-level model is an ensemble based on different other first-level models.

Denoised *soft labels* are then extracted from out-of-fold predictions of the ensemble: we take these predictions (in the form of $[0, 1]$ confidence values) as soft labels. Then, we split the training set into five folds again and train simple fully-connected models with different depth, width, and activation functions.

Finally, we make out-of-fold predictions again, take each model’s features on the penultimate layer (before final softmax), concatenate, and feed them into a new trainable classification layer. The feature extractors are frozen during training. This procedure yielded the final metric $\text{GAP} = 0.88729$.

Throughout this section, we refer to evaluation results summarized in Table 1, which shows whether the model uses frame-level features, its $\text{GAP}@20$ and BCE scores, and whether it became part of the final ensemble.

	Model	Fr.	GAP@20	BCE	Ens.
	Final ensemble	✓	0.88729	—	✓
1	ResNetLike + soft labels	✗	0.87417	9.2×10^{-4}	✓
2	ResNetLike + mixup	✗	0.86105	9.7×10^{-4}	✓
3	ResNetLike over linear combinations	✓	0.85325	1.02×10^{-3}	✓
4	ResNetLike + soft ranking loss	✗	0.85184	—	✓
5	AttentionNet	✓	0.85094	1.08×10^{-3}	✓
6	LSTM-Bi-Attention	✓	0.84645	1.04×10^{-3}	✓
7	Time Distributed Convolutions	✓	0.84144	1.0×10^{-3}	✓
8	VLAD-BOW + learnable power	✓	0.83959	1.1×10^{-3}	✓
9	Video only ResNetLike	✗	0.83212	1.1×10^{-3}	✓
10	Time Distributed Dense Sorting	✓	0.83136	—	✗
11	EarlyConcatLSTM	✓	0.82998	1.2×10^{-3}	✓
12	Time Distributed Dense Max Pooling	✓	0.82656	1.1×10^{-3}	✓
13	Self-attention (transformer encoder)	✓	0.8237	1.2×10^{-3}	✓
14	10000 clusters + ResNetLike	✓	0.7900	1.3×10^{-3}	✓
15	Audio only ResNetLike	✗	0.50676	2.5×10^{-3}	✓
16	Bottleneck 4 neurons	✗	0.41079	2.9×10^{-3}	✓

Table 1. Evaluation results for basic models, ordered from best to worst.

Video-level models. Our main first-level model is based on a ResNet-like architecture previously introduced for this problem in [5]; this architecture is shown on Fig. 3. The model takes three hyperparameters as input: inner_size, av_id_block_num, concat_id_block_num, and dropout. By default we used inner_size=2048, av_id_block_num=1, concat_id_block_num=1, dropout=0.5.

We have not found any strong correlation between network depth and final GAP value during our tests, yet our best first-level validation score GAP = 0.86105 was achieved by a ResNet-like model with parameters (4, 4, 0.4) and mixup $\alpha = 0.3$, trained with binary cross-entropy as the loss function (row 3 of Table 1).

Mixup One important idea that helped improve first level models was *mixup*, a recently developed new regularization technique [6]. Mixup has been empirically shown to deliver better accuracy on the validation set than original labels. The mixup method produces “virtual” training samples as linear combinations of existing training samples and their targets:

$$x = \lambda x_i + (1 - \lambda)x_j, \quad y = \lambda y_i + (1 - \lambda)y_j,$$

where (x_i, y_i) and (x_j, y_j) are feature-target vectors sampled from training data, and $\lambda \in [0, 1]$ is drawn from a beta distribution, $\lambda \sim \text{Beta}(\alpha, \alpha)$, where α is the main parameter of mixup; in our experiments we set the mixup α parameter to 0.4. Most of our first level models use mixup, and it improves results significantly; e.g., the basic ResNetLike architecture with mixup has GAP = 0.86105 (row 2 of Table 1) compared to GAP = 0.85846 without it.

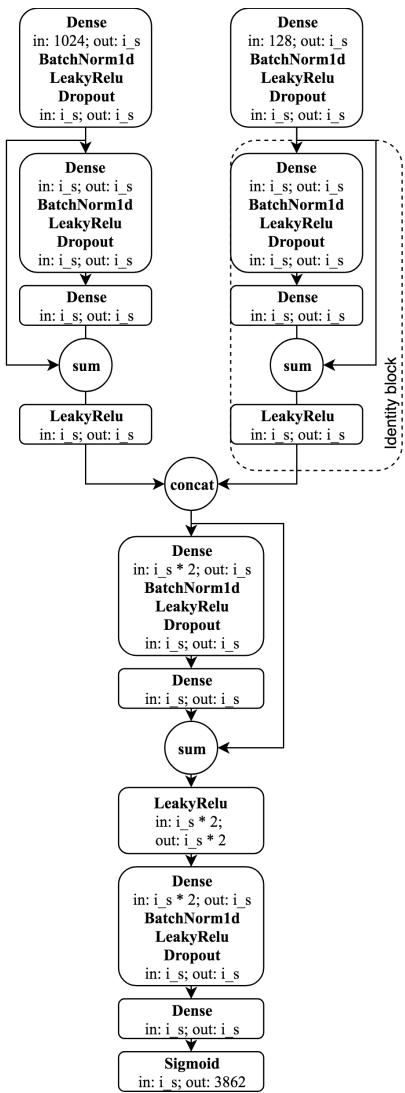
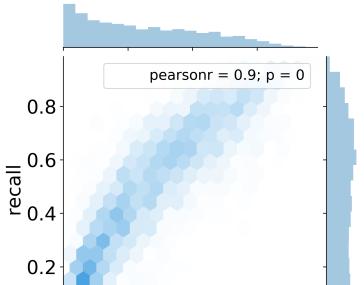
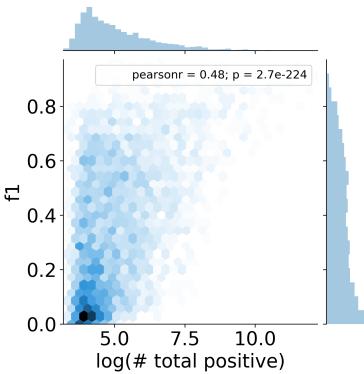


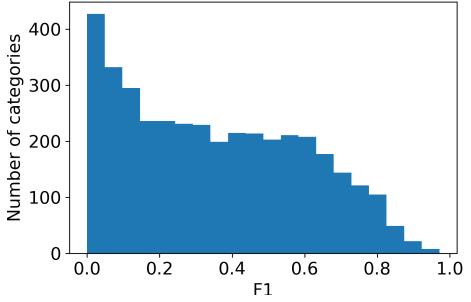
Fig. 3. ResNet-like architecture. i_s , inner size, denotes the size of hidden layers.



(a)



(b)



(c)

Fig. 4. Error analysis: (a) precision-recall heatmap; (b) F1-score as a function of the number of available training samples; (c) distribution of labels according to their final F1-score.

Recurrent frame-level models. Along with aggregated video-level features, the dataset also provides temporal frame-level representation of the videos. We have conducted multiple experiments with recurrent neural networks based on LSTM and GRU units to incorporate temporal features into the model. We have tried both unidirectional and bidirectional LSTM-based networks with 2 hidden layers size of 1024 followed by fully-connected layer size of 2048; we call this basic RNN model *EarlyConcatLSTM* because it receives as input concatenated audio and video features (row 11 of Table 1).

Another approach was to use a learnable bag-of-words representation. We used the *VLADBoW* model proposed in [7]. One major difference was to introduce a learnable power coefficient, that is, we learned it as

$$\text{BOW}(k) = \sum_{i=1}^N a_k(y_i) \text{ with } y_i = [Wx_i + b]_+^p \text{ and } a_k = \frac{e^{y_{ik}}}{\sum_{j=1}^k e^{y_{ij}}}$$

with trainable p , getting as a result $p = 0.628$. This model achieved GAP = 0.83959 (row 8 of Table 1).

Attention-based frame-level models. Recent research indicates that attention-based networks often outperform classical RNNs in similar tasks [8]. We have tried several attention-based architectures:

- (1) encoder architectures similar to the ones from the *Transformer* model [9] yielded GAP = 0.8237 (row 13 of Table 1); we used stacked multihead self-attention mechanism with and without position coding, and Transformer without position coding was much more successful; this may indicate that most labels could be detected by taking into account individual frames and disregarding relations between frames and their mutual temporal positions;
- (2) *stacked attention* model, which contains multiple layers, each computing a single attention vector for a sequence, concatenates them to obtain the feature vector at every time step, and then performs global average pooling, reaching a very high GAP = 0.85094 (row 5 of Table 1).

Time-distributed models. One more group of approaches is to use time-distributed dense layers with different poolings to finally convert them to a single vector; we have used three different approaches with time-distributed layers. Our simplest model uses a single fully connected layer and frame-wise max pooling afterwards. This approach yields GAP = 0.82656 (row 12 of Table 1). The second approach is to apply a ResNet-like model to every frame in a video with additional two dense layers on the cumulative output of every per-frame model sorted by confidence; this gave us GAP = 0.83136 (row 10 of Table 1). Finally, we tried time-distributed convolutional layers; the best convolutional model contains several layers of convolutions followed by max-pooling for video and audio separately, then concatenating the resulting features. This approach yields GAP = 0.84144 (row 7 of Table 1).

Loss function selection. As the main loss function we used the binary cross-entropy (BCE). However, since the primary objective was in terms of a ranking metric, some models were trained using the soft ranking loss; in our experiments, for equivalent models this loss has led to slightly worse GAP scores than BCE. Reweighting BCE using weights derived from mispredicted examples did not show any significant improvements in validation score.

The global average precision metric is based on the ranking among predictions rather than their specific scores. Thus, we applied batch-wise ranking in the following way. First we extract top 30 scores corresponding to negative labels for every sample in the batch. Then we extract all scores corresponding to positive labels in the batch. Finally, we apply the following pairwise ranking loss: $L(p_i, n_j) = \log(1 + \exp(n_j - p_i + 1))$, averaging it across all pairs of scores in the batch. This approach led to, e.g., GAP = 0.85184 (row 4 in Table 1) for ResNetLike compared to GAP = 0.85325 without the ranking loss (row 3), but ranking-based models were still useful for the final ensemble.

We also experimented with standard hinge ranking loss function, but found that hard thresholding is harmful for model convergence and the GAP metric. We have also tried the following related ideas with no strong positive effect (we believe that in this case it is important to report negative results as well):

- (1) penalizing predicted confidence for a label by cross-validation classification accuracy for that label (the smaller the accuracy, the stronger the penalty); the intuition here is that noisier labels must receive marginally less confident scores;
- (2) inducing noise by flipping labels at random (with higher probability for labels with smaller cross-validation accuracy);
- (3) loss function correction as proposed in [10];
- (4) dropout technique from [11];
- (5) test time dropout with averaging the predictions.

Final ensemble and distillation. Our general approach, as shown in the beginning of this section, implies a general ensemble model which is then used to obtain soft labels (similar to the knowledge distillation approach as shown in [12]) used instead of noisy ones to improve the final model’s quality.

We used an ensemble of 115 first level models to prepare the final model. On the first level, we used 95 video-level and 20 frame-level models; all of them are neural networks due to the origin and nature of the data provided. However, the ensemble itself was done with a gradient boosting model, namely LightGBM [13]. Distilled soft labels produced by this ensemble allowed us to achieve GAP = 0.88729 with the final compressed model.

4 Experimental Evaluation

In this section we present our experimental results on the YouTube-8M dataset.

Data preparation and augmentation. To make the resulting models more robust, we initially shuffled all data and split it into 5 cross-validation folds to train and validate all our models. Throughout the paper, we estimate the validation score as GAP measured on the hold-out fold, and the final GAP scores shown in Table 1 are calculated as an average of 5 folds.

The effectiveness of data augmentation in various tasks, especially related to computer vision, has been known and empirically validated for a long time [14–16]. Despite the fact that the initially provided features were preprocessed embeddings rather than images of frames, we performed several experiments with data augmentation techniques. In particular, we tested statistical bootstrap, using a weighted average for a random subset of frames and mixup [6]. Some models were trained on random subsets of available features.

Compact representations and single modality experiments. We have conducted experiments to see how performance (in terms of the validation metric) of the model degrades depending on the size of the inner representation within our model. Using only 4 neurons in hidden layers has allowed us to achieve GAP equal to 0.41079.

We have also trained separate audio-only and video-only video-level models based on a single modality only, achieving the best GAP scores of 0.83493 and 0.50676 respectively. These models were included into the final ensemble of first-level models for the sake of diversity.

Training parameters. All models in our comparison were trained and validated with multiple NVIDIA P40 GPUs. In our experiments, we have attempted to tune the mini-batch size and learning rate. The experiments have shown the following tradeoff: lower mini-batch size increases the convergence rate but the final quality of the model can deteriorate slightly. At the same time, lower batch sizes mean that an experiment takes up significantly less GPU memory (the memory footprint is basically linear in the batch size), which is an important argument for lower batch sizes in practice. Note that larger mini-batch sizes require warmup which does not let one use high learning rates right away [17].

Results. The main results of evaluating our models are summarized in Table 1. We have already discussed our models in detail in Section 3. One general result that can be seen from Table 1 is that, rather surprisingly, video-level models on the first level consistently outperform frame-level models despite the fact that they receive far less information about a given video. This probably means that there is still a lot of work to be done to improve recurrent and/or attention-based approaches to video understanding—after all, our final ensemble does use frame-level models and hence this information. Moreover, frame-level models have high importance in the gradient boosting model, which shows that they add a lot of diversity and new information to the ensemble.

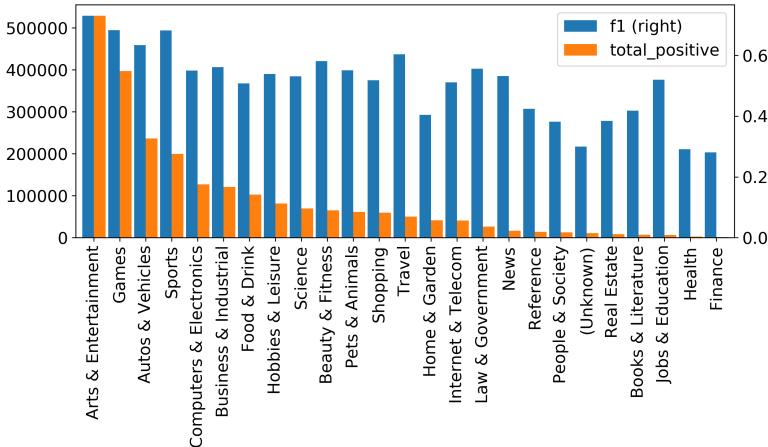


Fig. 5. Classification accuracy across verticals.

Error analysis To analyze the quality and characteristic features of our final model, we performed error analysis according to the following methodology. We take top 20 predicted labels for a video and classify these labels into true positive, false positive, and false negative classes. A label is true positive if it is positive and it has predicted score higher than all negative labels for that video; a label is false positive if it is negative and has predicted score higher than at least one positive label for that video; finally, a label is false negative if it is positive and has predicted score less than at least one negative label for that video.

Figure 4 shows the error analysis for our ResNet-like model distilled against the ensemble. Fig. 4a shows that false positives are marginally more probable than false negatives, but not by much, which shows that the model is well balanced. Fig. 4b illustrated an expected effect: the more training examples we have, the higher is the accuracy; note, however, that the final model can also accurately predict many rare labels. Nevertheless, Fig. 4c shows that vice versa, there are many categories that the model cannot predict accurately; these are mostly rare categories, but for some the reasons might be different and might require further study. Finally, Fig. 5 shows classification accuracy across verticals (coarse-grained labels) in comparison to the number of positive examples in the corresponding training subset (one fold). Again, we see that generally more popular categories are easier to predict, but the effect varies significantly.

5 Conclusion

In this work, we have presented our approach for large-scale video understanding implemented as part of the 2nd YouTube-8M Video Understanding Challenge. Our general idea is to use a complex model, a large ensemble of heterogeneous

neural networks, to produce soft labels for training set examples. The soft labels are then used as input for a much simpler model which is designed to fit into hardware requirements. We have shown the validity of our approach and presented a solution that took 4th place in the challenge.

The most promising direction for further work appears to be improving frame-level models: none of these models in our experiments outperformed video-level models (that have less information about the video), but they still provide important diversity for the ensemble, which shows that the information actually is there and can probably be extracted in a single first level model too. We also point out that the final model has very uneven quality across predicted labels; this probably means that in practice, different categories of labels may require different approaches.

References

1. Abu-El-Haija, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S.: Youtube-8m: A large-scale video classification benchmark. CoRR **abs/1609.08675** (2016)
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09. (2009)
3. Daubechies, I.: Ten Lectures on Wavelets. Volume 61 of CBMS-NSF Regional Conference Series in Applied Mathematics. (1992)
4. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? Journal of the ACM (JACM) **58**(3) (2011) 11
5. n01z3: Solution for google cloud & youtube-8m video understanding challenge. https://github.com/n01z3/kaggle_yt8m (2017)
6. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations. (2018)
7. Miech, A., Laptev, I., Sivic, J.: Learnable pooling with context gating for video classification. arXiv preprint arXiv:1706.06905 (2017)
8. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: Advances in Neural Information Processing Systems 30. Curran Associates, Inc. (2017) 5998–6008
9. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. (2017) 5998–6008
10. Natarajan, N., Dhillon, I.S., Ravikumar, P.K., Tewari, A.: Learning with noisy labels. In Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q., eds.: Advances in Neural Information Processing Systems 26. Curran Associates, Inc. (2013) 1196–1204
11. Jindal, I., Nokleby, M.S., Chen, X.: Learning deep networks from noisy labels with dropout regularization. CoRR **abs/1705.03419** (2017)
12. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
13. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: Lightgbm: A highly efficient gradient boosting decision tree. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: Advances in Neural Information Processing Systems 31. Curran Associates, Inc. (2018) 3147–3156

Advances in Neural Information Processing Systems 30. Curran Associates, Inc. (2017) 3146–3154

- 14. Ciresan, D.C., Meier, U., Gambardella, L.M., Schmidhuber, J.: Deep big simple neural nets excel on handwritten digit recognition. CoRR **abs/1003.0358** (2010)
- 15. Perez, L., Wang, J.: The effectiveness of data augmentation in image classification using deep learning. CoRR **abs/1712.04621** (2017)
- 16. Taylor, L., Nitschke, G.: Improving deep learning using generic data augmentation. CoRR **abs/1708.06020** (2017)
- 17. Goyal, P., Dollár, P., Girshick, R.B., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch SGD: training imagenet in 1 hour. CoRR **abs/1706.02677** (2017)