

Learning to Localize Temporal Events in Large-scale Video Data

Mikel Bober-Irizar
University of Cambridge
mikel@mxbi.net

Miha Skalic
University Pompeu Fabra

David Austin
Intel Corporation
david.j.austin@intel.com

Abstract

We address temporal localization of events in large-scale video data, in the context of the Youtube-8M Segments dataset. This emerging field within video recognition can enable applications to identify the precise time a specified event occurs in a video, which has broad implications for video search. To address this we present two separate approaches: (1) a gradient boosted decision tree model on a crafted dataset and (2) a combination of deep learning models based on frame-level data, video-level data, and a localization model. The combinations of these two approaches achieved 5th place in the 3rd Youtube-8M video recognition challenge.

1. Introduction

Understanding the contents of a scene from its video and audio features is a central capability of AI based systems today. There are many algorithms available to perform a classification task when presented with such input features. What is not as common but perhaps as valuable if not more so is to be able to take an long video and search for a given labeled scene. This scenario is much more common in the real world applications where video editing is costly and global data and meta-data can be noisy. Practical use cases for such kind of applications could include but not limited to: (1) police forces searching for "use of force" videos out of a large video library to properly asses and train their employees, (2) a security firm identifying hazards or threats from a repository of video, or (3) a video hosting service allowing users or advertisers to specifically identify scenes of interest of their users from a large database. Some recent approaches to localizing temporal events in a video include detecting audio-video events with attention[12] and using a modified R-CNN architecture to localize temporal context of actions[3].

Our contribution is to demonstrate that applying traditional machine learning and deep learning approaches to pre-extracted features is capable of achieving relatively high performance without the need for more expensive action

recognition networks or deep ensembles of many diverse networks.

2. Dataset & Competition Setup

The 3rd Youtube-8M video recognition challenge builds on an updated version of the original Youtube-8M dataset[1], known as the Youtube-8M segments dataset. This dataset contains 6.1 million YouTube videos, each labelled with some of 3862 classes (on average 3 per video). Due to dataset size constraints, instead of providing full audiovisual streams, each video is sampled at 1 frame per second (up to a maximum of 300 seconds), and each frame is passed though pre-trained Inception-v3[11] network as well as an acoustic model[5], followed by PCA and quantization to generate a pre-extracted 1152-D feature vector for each frame. This leads to a dense 2 dimensional matrix of features for each video in the dataset. The first two Youtube-8M competitions[7] were built upon this dataset, and so there is much existing research into solutions for the video-level classification problem.[8][9][10]

In addition to the above data, Youtube-8M segments adds 237,000 human-labelled *5 second segments* within videos - each labelled segment either contains or does not contain a specific class (out of a subset of 1000 classes). Fewer than 50,000 (1%) of the videos in the dataset are labelled in this way, with each of these videos containing on average 5 labelled segments. This yields a segment-level labelling which is **extremely sparse**, with only 0.0001% of all possible segment-class pairs being labelled, and fewer than 250 total segment labels per class.

The goal of the 3rd Youtube-8M video recognition challenge is to predict *segment-level* class labels for each video in the test set. Given the scarcity of segment-level training examples, it is vital for any solution to leverage the vast number of video-level training examples in some way during training - the problem is one of weakly supervised learning. In essence, we have to learn to localise classes within videos without knowing (or rarely knowing) where the classes appear in the training set examples.

3. Evaluation

The Kaggle competition is evaluated using Mean Average Precision (MAP@K) where $K = 100,000$. The Average Precision (eq. 1) is computed on a *per class* basis, and averaged across the 1000 classes in the dataset.

$$MAP@K = \frac{\sum_{k=1}^K P(k) \times rel(k)}{N_{pos}}, \quad (1)$$

where $P(k)$ is the precision in the first k predictions, $rel(k)$ is 1 if prediction k is correct and 0 otherwise, and N_{pos} is the total number of positive labelled segments for that class.

Thus for each class, the task is to predict the top 100,000 labelled video segments which may contain that class in order of likelihood. As the segment labelling is also extremely sparse in the test set (and we do not know which segments are labelled), it is necessary to predict many more than 100,000 segments for each class in order for 100,000 predictions to be ‘considered’. In practice, the limiting factor on number of predicted segments is the timeout on Kaggle’s evaluation servers (which will fail if a submission takes more than a few minutes to score) rather than the limit of $K = 100000$.

4. Video-level Baseline

Because the evaluation metric is evaluated on a per-class basis, and not per-video (for example, ranking segments within videos), this means that it is possible to achieve a high score simply by calculating a likelihood for each (video, class) pair - as in previous competitions - and assuming that within each video, all segments have the same probability of containing that class. This yields a pseudo-ranking for the segments which can be submitted. We took the winning model from the 2018 Youtube-8M competition[10] (with a GAP score above 0.890), and used the naive approach detailed to achieve a score of **0.68182 MAP**.

Thus, it is important to note that a large part of the evaluation metric for this dataset still relies on ordering videos based on likelihood of each label, in addition to ranking segments within each video, and a good solution must be able to effectively do both.

To this end, here we present and contrast two machine learning approaches that have been developed to localize target events in videos. The first approach utilizes gradient boosting decision trees, while the second approach is a more ‘conventional’ deep learning-based approach. These two methods are complimentary, trained in parallel, and are subsequently combined to produce our final approach.

5. Gradient Boosting Trees for Event Localization

5.1. Frame-Level Model

First, to leverage the large number of *video-level* labels available (7 million training examples), a basic feed-forward neural network was trained to predict labels based on the provided feature vectors averaged across the time dimension to form a single 1152-D feature vector per video. This network was trained on video-level average data (with the available video-level labels), but is then used to predict on **individual frames** instead. We refer to a prediction of likelihoods from this model on frame $X \in \mathbb{R}^{1152}$ as $fl(X) \in \mathbb{R}^{1000}$.

5.2. Frame-level dataset construction

Next, for each of the 1000 classes: a dataset of frames was constructed by taking all the labelled 5-second segments for that class, and sampling 9 frames from each segment (two on either side of the labelled five, assuming they also have the same class), and concatenating these segments to form a frame-level dataset for each class. This approach yielded approximately 1000 “definitely positive” and 1000 “definitely negative” frames for each class (frames from segments that were human-labelled as having or not having the class). As this is a relatively small dataset, several thousand additional negative frames were randomly sampled from the large training set from videos which are not labelled with the class in question.

It is important to note that while we could have sampled additional negatives (each class has approx. 1 billion “probably negative” frames in the entire dataset), this was not done as these frames tend to be trivial for a model to classify, distracting from the real problem of distinguishing between positive and negative segments *within the same video* - which was found experimentally to be much more difficult for a model to do.

Each frame descriptor from this dataset X is then passed through the trained video-level model (as if it was a video, to get a 1000-class prediction vector) and the results are concatenated to form a new feature vector $\{X, fl(X)\} \in \mathbb{R}^{2252}$. This forms our frame-level dataset.

5.3. Gradient Boosted Trees Model

After this frame-level dataset has been constructed (around 5000×2252), for each class c a set of XGBoost[4] models xgb_c is trained with a 5-fold cross-validation approach, for a total of 5000 models (each model takes a few seconds to train - about 24 hours total). Because each class only has on average 125 positively labelled training segments from which the entire training set is derived, there is a large variance between folds which use different segments

for validation. Hence, at test time, the average of these 5 models for each class are taken to improve performance.

5.4. Model composition

The XGBoost model xgb_c has been trained on a specially crafted (and thus **biased**) subset of the full set of frames. This is because the Youtube-8M segments data only has segments labelled for videos **where the video itself has the class c** - the training data is very heavily biased towards these videos.

One solution to this problem would be to construct a frame dataset which accurately reflects the distribution of frames or 1 in the whole dataset, this however presents two challenges:

1. Very few frames in the dataset are labelled (on a segment level), so creating a representative data would be difficult - it is not as simple as just subsampling the data.
2. This has the risk of transforming the problem into one of mostly detecting which videos have the class versus which videos are completely unrelated, instead of the more difficult (and important) task of distinguishing positive and negative frames *within videos that are labelled with the class*.

Instead, we model the output of xgb_c as a conditional probability:

$$xgb_c(\{X, fl(X)\}) = P(X \in c | V_X \in c) \quad (2)$$

where V_X is the video of frame X . Essentially, our XGBoost model predicts the probability that a frame X has class c , given that the video the frame was from has class c .

Our goal is to calculate the probability $P(X \in c)$. We can use Bayes' rule on (Eq. 2) to get

$$P(X \in c | V_X \in c) = \frac{P(V_X \in c | X \in c) \cdot P(X \in c)}{P(V_X \in c)} \quad (3)$$

$P(V_X \in c | X \in c)$ is 1, as a video of segment in class c must necessarily also be in class c . And so:

$$P(X \in c | V_X \in c) = \frac{P(X \in c)}{P(V_X \in c)} \quad (4)$$

$$P(X \in c) = P(X \in c | V_X \in c) \cdot P(V_X \in c)$$

$$P(X \in c) = xgb_c(\{X, fl(X)\}) \cdot P(V_X \in c)$$

Now we just need the video-level probability $P(V_X \in c)$. Luckily, we can use the state-of-the art video-level model[10] from Section 4 to estimate a video-level probability $vl_c(V_X) = P(V_X \in c)$. The final prediction for

each frame is the product of this video-level model and the XGBoost model - and the prediction for each 5 second segment is this probability averaged over the 5 frames in the segment.

$$P(S \in c) = \frac{1}{5} \sum_{i=1}^5 xgb_c(\{X_i, fl(X_i)\}) \cdot vl_c(V_S) \quad (5)$$

where S is a segment containing frames $X_1 \dots X_5$. By averaging the likelihoods of the 5 frames in each segment, we can consider more information for our prediction instead of just using information from a single frame.

5.5. Model weighting

The product of the two models, in a sense, 'equally weights' the class probability predictions estimated by our XGBoost model and the video-level model. Where both models are perfect predictors on their respective datasets, this works fine (as the equations show) - however, the reality is that some models are more effective than others, and so it may be useful to change these weightings.

The product can also be interpreted as a geometric mean ensemble between the two models (as the geometric mean is just the product followed by the square root, which would not affect the ranking of segments for MAP). Thus, we can introduce an additional parameter p to control the weighting between the models

$$P(X \in c) = (xgb_c(\dots))^p \cdot vl_c(V_S) \quad (6)$$

A $p > 1$ provides higher weight to XGBoost, while $p < 1$ provides higher weight to the video-level model. A higher p in general produces smaller (skewed) probabilities to be predicted, but as MAP only cares about *ranking* of segments, there is no need to correct for this skew. This can be seen as analagous to weighting models in a linear ensemble: $p \cdot (xgb_c(\dots)) + vl_c(V_S)$.

Experimentally, $p = 4$ was found to be a good value, improving the MAP of this model from 0.77552 to **0.79545**.

6. Deep learning for event localization

6.1. 5-frame models

As the goal is to label 5 second segments, we built models that process 5 frames at a time (a single segment) without considering other information from elsewhere in the video.

In particular, 2 models were trained based on the previously successful NetVLAD[2, 8] network architecture and one based on a deep bag of frames (DBoF)[1] network architecture. The models were first pre-trained on the 2nd Youtube-8M dataset, by selecting 5 frames from the videos at random, and assuming that these segments inherit the

classes of the video - yielding a large but very approximate dataset. Finetuning on specifically annotated 5 second sequences followed. In the finetuning process any unlabelled classes (neither positively or negatively, due to the extremely sparse labelling) were masked out and not trained on, to avoid affecting the model learning process with incorrect labels. Furthermore, initial pretraining on the bigger dataset was carried out for longer ($\approx 50K$ and $\approx 10K$ steps at a batch size of 1024 and 2048 for the VLAD and DBoF models respectively) than finetuning (≈ 500 iterations at previously mentioned batch sizes). This was done to avoid overfitting to the comparatively tiny segment-annotated dataset.

Single VLAD model achieved MAP of ≈ 0.74 and combining the three models improved performance further. Empirically we observed that multiplying probabilities (geometric mean) gave a better MAP uplift than averaging (arithmetic mean) the model predictions (Table 1).

6.2. Global model

A 5 second snippet does not give all the possible useful information about a video - context is also important; for example, if a video is all about cars, then it is more likely that a certain given segment is also about cars. For this reason, we want to leverage ‘global’ information from the video, where just relying on 5 frames and ignoring broader video context might result in misclassification and lead to suboptimal performance.

To tackle this, we combined the segment-level model outputs with label predictions for the whole video. This was done in a similar way to the Section 5.4; a mean was taken between the local predictions of the 5-frame models and the . It was found that the geometric mean significantly outperformed the arithmetic mean for combining these models (Table 1).

6.3. Localization model

The model described in previous section helped with overall predictive performance. However, the context modelling video-level model did not add any temporal differentiation. Hence, we introduced a localization model to help prioritize sections of video that contain the target classes.

The localization model was implemented as follows. The target class is passed through an embedding layer converting our sparse one hot encoding of 1000 classes to a low 32-dimensional continuous representation. This representation is then concatenated with the 1152-dimensional frame-level representation, before being passed to an dynamic length bidirectional long short-term memory (LSTM)[6] network. Next, the LSTM output is processed by a fully-connected layer to yield a single logit per frame. Finally, the sigmoid activation is applied to yield class probabilities and this is trained with backpropagation. Importantly, unan-

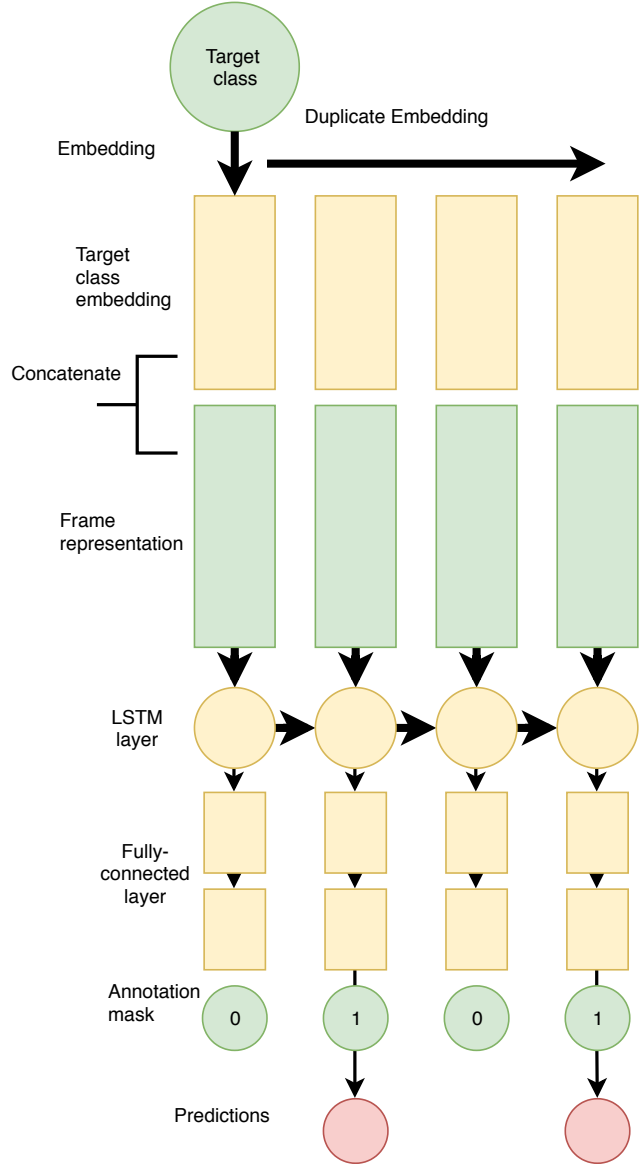


Figure 1. Schematic representation of the localization network. Inputs are colored green, while latent values and outputs are colored yellow and red, respectively.

notated segments were masked out of the loss calculation not to interfere with model training via backpropagation. Schematic representation of this network is shown in Figure 1. During inference all videos were passed through the network in combination with all 1000 target classes in order to get class specific probabilities.

6.4. Combining the models

Final score ($S_{v,s,c}$) for a segment s of video v for class event c was calculated as follows:

$$S_{v,s,c} = p_{video} \cdot p_{loc} \cdot (p_{VLAD1} \cdot p_{VLAD2} \cdot p_{DBoF})^{2/3}. \quad (7)$$

Models	MAP score
VLAD Model (5-frame)	0.74056
3x 5-frame (arithmetic mean)	0.76124
3x 5-frame (geometric mean)	0.76523
3x 5-frame, global model	0.78209
3x 5-frame, global model, localization model	0.78732

Table 1. MAP score for different model or model combination predictions. MAP scores are reported based on the private Kaggle leaderboard.

In the equation p_{video} are predictions for video, p_{loc} are localization model predictions and 5-frame models predictions are represented as p_{DBoF} and p_{VLAD} . Video model gives class predictions for each video, while 5-frame models and localization model give class predictions for each segment in video. For each target class the segments were sorted by final score in descending order and top ranking segments were reported.

We would like to note that due to time and computational constraints the models were not (exhaustively) tuned. Rather, they serve as a proof of concept and with proper tuning the performance can be further improved.

7. Final meta-model

Model	MAP score
Video-Level Baseline (Sec. 4)	0.68182
Gradient Boosting Approach (Sec. 5)	0.79545
Deep-learning approach (Sec. 6)	0.78732
Rank Average	0.80459

Table 2. MAP score for different model or model combination predictions. MAP is reported based on the Kaggle private leaderboard score.

Predictions from the models described in previous two sections were, ensembled into final predictions. This was done using a rank average: for each class, the predictions from the two models are assigned a score based on their rank in the prediction list, this rank is averaged, and then the list is sorted on the average rank in all the models. This allows the models to be ensembled without having to keep around the full table of probabilities, and is robust to when the probability distributions of the two models are very different (as is likely the case here). The ensembled prediction achieved MAP score of **0.80459** (Table 2) and ranked 5th as part of the 3rd Youtube-8M video recognition challenge. Code to reproduce the results presented here, including the chosen hyperparameters, is available at <https://github.com/mxbi/youtube8m-2019>.

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] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pfister, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016.
- [3] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster R-CNN architecture for temporal action localization. *CoRR*, abs/1804.07667, 2018.
- [4] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. *CoRR*, abs/1603.02754, 2016.
- [5] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [6] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [7] Joonseok Lee, Apostol Natsev, Walter Reade, Rahul Sukthankar, and George Toderici. *The 2nd YouTube-8M Large-Scale Video Understanding Challenge: Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, pages 193–205. 01 2019.
- [8] Antoine Miech, Ivan Laptev, and Josef Sivic. Learnable pooling with context gating for video classification. *CoRR*, abs/1706.06905, 2017.
- [9] Eng-Jon Ong, Sameed Husain, Mikel Bober-Irizar, and Mirosław Bober. Deep architectures and ensembles for semantic video classification. *CoRR*, abs/1807.01026, 2018.
- [10] Miha Skalic and David Austin. Building a size constrained predictive model for video classification. September 2018.
- [11] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567, 2015.
- [12] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. *CoRR*, abs/1803.08842, 2018.