

AUDIOVISUAL HIGHLIGHT DETECTION IN VIDEOS

Karel Mundnich^{**} Alexandra Fenster[‡] Aparna Khare[‡] Shiva Sundaram[‡]

^{*}Signal Analysis and Interpretation Lab, University of Southern California, Los Angeles, CA 90089

[‡]Amazon Lab126, Sunnyvale, CA 90489

ABSTRACT

In this paper, we test the hypothesis that interesting events in unstructured videos are inherently audiovisual. We combine deep image representations for object recognition and scene understanding with representations from an audiovisual affect recognition model. To this set, we include content agnostic audio-visual synchrony representations and mel-frequency cepstral coefficients to capture other intrinsic properties of audio. These features are used in a modular supervised model. We present results from two experiments: efficacy study of single features on the task, and an ablation study where we leave one feature out at a time. For the video summarization task, our results indicate that the visual features carry most information, and including audiovisual features improves over visual-only information. To better study the task of highlight detection, we run a pilot experiment with highlights annotations for a small subset of video clips and fine-tune our best model on it. Results indicate that we can transfer knowledge from the video summarization task to a model trained specifically for the task of highlight detection.

Index Terms— audiovisual, highlight detection, video

1. INTRODUCTION

Finding highlights or *interesting* segments in audiovisual content has become an increasingly relevant task in video processing. With the proliferation of smartphones and widely accessible video cameras, the amount of audiovisual content and the demand to share them has increased over the years. This has also increased the need to develop automated tools to find the most interesting segments while filtering out portions that are redundant. This is a non-trivial task in an unconstrained setting as finding relevant portions (rather, what we define as relevant or interesting) of audiovisual content depends on the context, nature of the videos (i.e. moving or static, egocentric point of view vs. not [1]) and topic or genre.

To this end, the problem of finding highlights in audiovisual content has been studied in different settings. For example, Wrede et al. [2] define hotspots in meetings as “regions of meetings that are marked by high participant involvement, as judged by human annotators”, while Lai et al. in [3] define summarization hotspots as “regions chosen for inclusion in a summary by human annotators”. These definitions capture the intrinsic subjectivity of the task: both are defined by the judgment of human annotators. These definitions also partially overlap with the definition of video summarization: “[...] to generate a short summary of the content of a longer video document by selecting and presenting the most informative or interesting materials for potential users.” as defined by Ngo et al. in [4]. However, the goals of both tasks are different: for highlight detection, we try to find those segments of audiovisual content that are most *interesting* (i.e., the best moments), while the goal for video summarization is to generate an overview of the content [5, 6].

1.1. Related work

A plethora of works have addressed both problems of video summarization and highlight detection. One of the first works in video summarization is by Smith et al. in [7], where the authors combine video features and language content to create video skims. In recent years however, visual features have been the focus for video summarization. For example, in [1], Gygli et al. use aesthetics/quality in frames, presence of landmarks, and face/person detection as features for summarization. Other works use visual features only [8, 9, 10]. A few have also used more than one modality. For example, Song et al. in [11] use the lexical information in video titles from YouTube to aid the summarization, while Jiang et al. [12] use video features as well as raw audio features for the same task.

For highlight detection, both visual-only and multimodal approaches exist. In [13], Yang et al. use visual 3D convolutional deep features and segments from edited videos to train an auto-encoder for highlight detection. Molino et al. in [5] use GIF generation from <http://www.gifs.com> as labels to find personalized highlights in videos using only visual features. Both of these only use visual features from videos. In terms of multimodal processing, [14, 15, 16] have presented novel approaches to finding highlights in sport videos, while [17] uses a multimodal approach for highlight detection in movies. These domains are however, constrained. In work not involving video, Makhervaks et al. in [18] propose using acoustic information from speech, lexical content from transcripts, and interaction statistics of participants to find hotspots in meetings, and found that the lexical content was the most informative modality.

1.2. Contributions

In contrast with the research discussed above, we take a modular supervised approach to the problem of highlight detection in unstructured and unconstrained audiovisual content. We combine representations from semantically meaningful models that take into account scene understanding, object recognition and face detection in vision, affect (arousal and valence states of individuals in the video) and synchrony in audiovisual content. We include audio-only features in the form of mel-frequency cepstral coefficients (MFCCs) that are popular in audio information retrieval problems. We study the performance of these modalities by training models over video summarization labels, whose predictions we use to create rankings for highlight detection. Moreover, we use these pre-trained models on in-house collected annotations over a small set of videos with social interactions, and fine-tune over this set of labels. We do so to explore the use of transfer learning from the video summarization task on a large unstructured dataset to the highlight detection task on a more specific domain as this technique has been successfully applied in both computer vision and speech. These visual, audiovisual and audio only representations stem from our motivation to analyze the relevance of audiovisual cues in summarization and video highlight detection tasks. While a truly end-to-end solution is our north star, we limit the scope of this study to analyze the performance of

^{*}Work performed during an internship at Amazon Lab126.

these modalities for our task.

2. METHODOLOGY

2.1. Dataset

We use the CoView dataset [19] that consists of 1500 videos. Each video in the train set contains summarization labels for each 5s-long segment, collected from 20 annotators, in the interval $[0, 2]$ (where 0 is not important at all and 2 is very important). The train partition consists of 1200 annotated videos, where the minimum video length is 100s. To the best of our knowledge, the method to aggregate the annotations as well as the inter-rater agreement are not reported. We only employ the train set, since there are no publicly available annotations for the test partition.

2.2. Extracted Features

We extracted the following features from the CoView dataset.

GoogLeNet To have frame-wise representations of objects in images, we extract 1024-dimensional embeddings from a pre-trained GoogLeNet [20] model.

Places365 To represent scene understanding for unstructured videos, we obtain 2048-dimensional embeddings from a ResNet50 [21] model pre-trained on the Places2 database [22]. We posit that videos with scenic views would be of high interest according to human judgment, which motivated including these embeddings into the model.

Affect We extract 128-dimensional embeddings for arousal and valence from the videos using the multimodal model proposed in [23], considering that portions of video with higher demonstrations of affect will be included in highlights.

Multisensory class activation maps (CAMs) (Owens et al. [24], Zhou et al. [25]) We obtain 14×14 foreground CAMs per frame for each video, following the same procedure as [24] to obtain CAM visualizations but without up-sampling the actual frame size. We assume that interesting portions of videos contain actions where audio and video are correlated.

Faces Following Gygli et al. in [1], we use a face detection module to detect faces in videos. As a feature, we use the area of the face detected with highest confidence (and 0 area in frames with no face detected). We assume that interesting portions of videos are likely to have people in them.

MFCCs We obtain 13-dimensional MFCCs from the audio sub-sampled to 16kHz, with a window length of 25ms and a window step of (0.25/29.97)s such that we obtain 4 samples per video frame. The audio features could be helpful in identifying attributes such as presence of speech, music and other acoustic sources [26].

We converted all videos to 29.97 frames per second prior to feature extraction.¹ After extraction, all features are then upsampled to 30 fps before training the models which allows the feature frames to be aligned with the CoView labels. The input to the model is all features corresponding to a 5 second labeled video segment.

2.3. Model

We propose a modular supervised model for the prediction of summarization labels. Fig. 1 shows the architecture: we employ different pre-processing schemes for different modalities. We then concatenate all the intermediate representations to get a single embedding vector before a bidirectional LSTM layer. We now describe each pre-processing module:

GoogLeNet, Places365 and Affect embeddings We use a dense layer for each embedding vector to reduce their dimension to 64.

Multisensory CAMs We employ 64 convolutional filters of size 14×14 per frame, whose output is passed to a dense layer for a

¹The 29.97 fps for feature extraction originates from the multisensory [24] model, which expects that sampling rate.

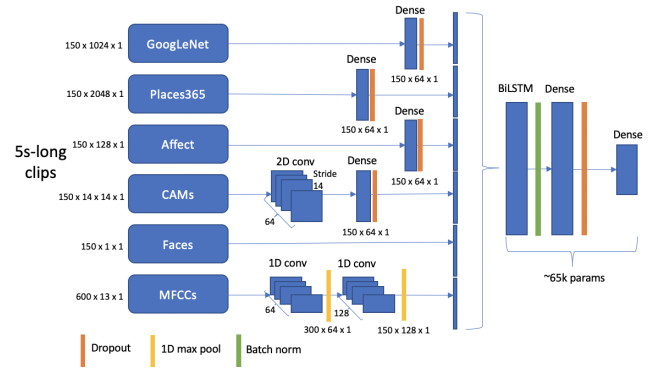


Fig. 1: Supervised model architecture.

final dimension of 64.

Faces Representation is directly concatenated with the other embeddings.

MFCCs We employ 2 convolutional layers (in time). The first one with 64 filters and the second one with 128 filters. Each convolutional layer is followed by a 1-dimensional, stride-2 max pool, for a final dimension of 128.

We use a dropout of 0.5 after each dense layer, and both recurrent and input dropout in the bidirectional LSTM layer. The LSTM and dense layers have 20 hidden units each, and the complete model has 314k params without considering the networks from which features have been extracted.

2.4. Loss functions

L1 or L2 loss functions only penalize for the point-wise differences between the model output and the labels, without considering relative scores between different segments. Therefore, we train with two loss functions: Canonical Correlation Coefficient (CCC) loss, and ranking loss. Furthermore, our experiments with L2 loss also showed inferior performance when compared with CCC loss.

Each training example to the model represents a 5-second long clip, with the corresponding summarization score from the CoView dataset.

We use the CCC loss because it penalizes both different ordinal variations in the sequences as well as differences in their means. Given inputs x , the labels y and the output from the model $f(x)$, CCC is defined as:

$$\rho_c(y, f(x)) = \frac{2\sigma_y\sigma_{f(x)}}{\sigma_y^2 + \sigma_{f(x)}^2 + (\mu_y - \mu_{f(x)})^2} \in [0, 1], \quad (1)$$

where μ and σ are the mean and standard deviation for each sequence. To convert this correlation to a loss, we use $\ell(y, f(x)) = 1 - \rho_c(y, f(x))$ as the loss.

Ranking loss allows the model to learn how to assign higher scores to more interesting clips from the CoView labels perspective. To train with margin ranking loss as defined by Yao et al. in [27], we sample the set \mathcal{P} of clips pairs (h_i, l_i) such that the CoView label value for one clip in the pair is strictly higher than the other clip's label value. Across all pairs, the margin ranking loss is defined as:

$$\sum_{(h_i, l_i) \in \mathcal{P}} \max(0, 1 - f(h_i) + f(l_i)), \quad (2)$$

where $f(\cdot)$ denotes the output score of the model.

2.5. Training

We train the network using 5s-long segments. We use Adam to train, with a learning rate of 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\varepsilon = 1 \times 10^{-7}$. For training with CCC loss, we use a batch size of 300, and train over 20 epochs, after which we pick the best performing model over the validation set, over all epochs. For training with ranking

loss, we use a batch size of 64, and train over 10 epochs, and the best model is selected the same way.

2.6. Evaluation metrics

We report the mean of NDCG and Precision at k scores across all videos in the dataset.

2.6.1. NDCG

Since we are computing rankings, we evaluate our methods using the Normalized Discriminative Cumulative Gain (NDCG) as described by Järvelin et al. in [28], which yields a value related to the quality of an estimated ranking. If we assume that each ranked item has an estimated relevance score \hat{r}_i (output of the model), then the Discriminative Cumulative Gain (DCG) over the top- k items of the estimated ranking is defined as:

$$DCG_k = \sum_{i=1}^k \frac{\hat{r}_i}{\log_2(i+1)}, \quad (3)$$

where the \log_2 -based denominator is a penalization term for the position of item i in the ranking. If we also know the ground truth relevances r_i , we can compute the *ideal* DCG (IDCG) by replacing the \hat{r}_i with r_i in Eq. 3. Then, the Normalized DCG is simply:

$$NDCG_k = \frac{DCG_k}{IDCG_k} \in [0, 1], \quad (4)$$

where a value of 1 implies perfect ranking.

2.6.2. Precision at k ($P@k$)

We use the Precision at k ($P@k$) score to evaluate how many of the top- k highlights we have been able to detect, based on the ranking method described above. The precision score is defined as ratio of true positives over all positives (TP: True positive, FP: False positive):

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (5)$$

or the number of selected items that are relevant. To compute $P@k$, we sort the labels and predictions in descending order, and threshold at the top- k .

3. EXPERIMENTS

We run three different studies to understand the contribution of different modalities, as well as the ability to transfer the learning from unstructured videos to videos within a given topic.

3.1. Experiments on CoView data using summarization labels

For both experiments, we run 5-fold cross-validation for each single feature or set of features and report cross-validation performance.

3.1.1. Relevance of single features

To understand the contribution of single features, we train the model from Fig. 1 with only one feature at a time. Since the sizes of the embeddings are different for some features, we keep the number of parameters of the layers after concatenation close to 15k. For MFCCs, we reduce the number of LSTM and hidden layer units from 20 to 13, and for faces, from 20 to 4.

3.1.2. Ablation experiments

To understand the impact of each feature in the full model, we run an ablation study, where we leave one feature out at a time from the model. Again, since the number of trainable parameters change, we adjust the number of hidden units to have approximately 65k parameters for each experiment.

We run ablation experiments with both CCC and ranking loss functions. Following Yao et al. [27], for experiments with ranking loss we build a siamese architecture based on the model described above, where both sub-networks are identical and share all parameters. For evaluation, we run every clip through a single sub-network to produce one score per clip. Further evaluation process is identical to CCC evaluation.

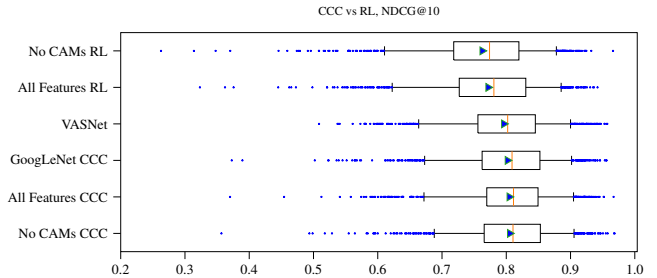


Fig. 2: Box plots for NDCG@10 for best performing models. Outer whiskers set to 5th and 95th percentiles. Triangles are at mean values, vertical lines are at medians and stars are outliers.

3.1.3. VASNet experiment

We compare our model with the VASNet model architecture proposed by Fajtl, Jiri et al. in [10]. Since this model was trained on frame-wise labels, and CoView only provides one label for a 5-second segment, we used the same score for all frames in the 5-second window for training. During inference, we computed the score of the 5-second segment as the mean of the score for each frame in the segment. VASNet was trained on GoogLeNet embeddings [20], so we find it directly comparable to our model trained on the same single feature (GoogLeNet column in Table 1).

3.2. Experiments on CoView with highlight labels

3.2.1. Annotations for highlight detection

For highlight annotations, we took a subset of 30 videos of the train dataset and collected annotations for the same 5s-long segments. We preferred lightly edited videos containing social interactions. Each video was annotated by three annotators. The annotators were instructed to watch the videos from beginning to end, and to annotate each 5s-long segment using the following instructions: “Please rate each clip using the scale below to label the segments of the video that contain highlights, defined as an outstanding part of a video over time. Note that a video may have no highlights.”, where the scale was: (1) *Definitely not a highlight*, (2) *Not a highlight*, (3) *Maybe a highlight*, (4) *One of the highlights*, (5) *Definitely a highlight*. As preprocessing of the highlight annotations, we scale the annotated values between 1 and 5 to the interval $[0, 2]$.

3.2.2. Fine-tuning

To fine-tune the model, we train the full model on a single fold with CoView labels. We then fine-tune on the 30 videos with highlight scores using 6-fold cross validation, with 25 videos for training and 5 for evaluation in each fold, with a batch size of 64 for 40 epochs, with a learning rate of 1×10^{-4} . The annotations obtained from the annotators are averaged to get the labels used for training.

4. RESULTS AND ANALYSIS

4.1. Experiments on CoView data and labels

4.1.1. Relevance of single features

We present the results using single features in Table 1, where the “Chance” column shows the chance values for all metrics. The results shown are the performance over 5-fold cross validation. For this task, we observe the best performance from the GoogLeNet embeddings (trained on ImageNet) across all metrics, followed by the Places365 embeddings (trained on Places2), while the least informative features are the proportion of the faces in the frames. For $P@5$, faces, affect, and multisensory CAMs are below chance.

4.1.2. Ablation experiments

For training with CCC loss, we present the results using different subsets of features in Table 2, where the “All CCC” column shows the results when training with all features. For this task, we observe that including most features in the model beats chance, and

Table 1: Single feature NDCG and precision with CCC loss. Highest scores in **bold**, lowest scores in *italics*. VASNet [10] for comparison.

Metric	k	Features					Chance	VASNet	
		GoogLeNet	Places365	Affect	MFCCs	CAMs			Faces
NDCG@ <i>k</i>	5	0.763	0.758	0.743	0.723	0.726	0.720	0.717	0.757
	10	0.802	0.799	0.783	0.765	0.769	0.763	0.756	0.796
P@ <i>k</i>	5	0.170	0.163	0.127	0.135	0.128	0.126	0.129	0.165
	10	0.309	0.306	0.258	0.262	0.257	0.260	0.251	0.293

Table 2: Leave-one-feature-out NDCG and precision using CCC loss with best models trained with Ranking Loss (RL) for comparison.

Metric	k	Features CCC					All CCC	All RL	RL	RL	
		No GoogLeNet	No Places365	No Affect	No MFCCs	No CAMs					No Faces
NDCG@ <i>k</i>	5	0.764	0.764	0.766	0.763	0.766	0.765	0.765	0.729	0.723	0.715
	10	0.804	0.802	0.804	0.804	0.806	0.804	0.805	0.772	0.765	0.763
P@ <i>k</i>	5	0.169	0.177	0.178	0.171	0.174	0.174	0.170	0.140	0.130	0.120
	10	0.312	0.307	0.307	0.312	0.316	0.311	0.312	0.265	0.256	0.263

removing multisensory CAMs produces the best cross validation results when we train with CCC loss. We also observe that except for the model without multisensory CAMs, most models have a performance similar to using GoogLeNet embeddings alone. We show the NDCG@10 box-plots for the best performing models in Fig. 2. The figure shows that although the mean and median NDCG values of the model trained with GoogLeNet alone are similar to the proposed model, the first quartile is higher for our proposed model.

We also observe that results from training with CCC loss are always superior to results from training with ranking loss for all ablation experiments performed. Fig. 2 reflects this superiority for best performing models within each loss as well. Best model trained with ranking loss is the one trained on all features.

4.1.3. VASNet experiment

First and third quartile, mean and median values for VASNet [10] model are shifted to the left in Fig. 2 in comparison with models trained with CCC loss, which indicates that all three CCC models improve over VASNet [10] results by a noticeable margin.

4.2. Experiments on CoView with highlight labels

4.2.1. Highlight Annotations

Fig. 3 shows the inter-rater agreement (Cronbach’s α) among annotators for all videos, with a mean of 0.524, max of 0.9, and min of 0.01. This figure shows the subjectivity of the task: even though there are videos for which the annotators agreed on the highlights, there are a large number of videos for which the agreement is low. We also show Kendall’s rank correlation with CoView labels, with a mean of 0.244, max value of 0.640, and min value of -0.230. This shows the difference in the tasks, and invites the question of whether we can transfer the knowledge from one set of labels to the other.

4.2.2. Fine-tuning results

We present the fine-tuning results in Table 3. The baseline model is trained on 30 videos. For fine-tuning, we initialize the model weights learned by training the model on summarization labels with all the features (All CCC in Table 2). The results show that fine-tuning can effectively transfer knowledge from the summarization task to the highlight task, with the fine-tuned model demonstrating an improvement over all the metrics. Both models, however, have a lower than chance NDCG score. NDCG depends not just on the relative ranking between segments, but also on the actual scores assigned to each segment. The precision values show that the transfer learned model ranks the segment better, but the NDCG metric show that the scores are not calibrated well. We will focus on this in our future work.

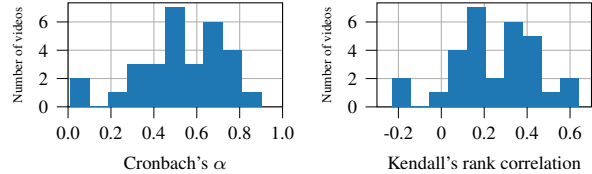


Fig. 3: Annotator agreements and correlations with CoView labels. (Left) Histogram with Cronbach’s α for the 30 annotated videos. (Right) Histogram with Kendall’s rank correlation between aggregated annotations and CoView labels.

Table 3: Fine-tuning results on highlight labels.

Metric	k	Chance	Baseline	Fine tuning
NDCG@ <i>k</i>	5	0.677	0.581	0.617
	10	0.721	0.612	0.665
P@ <i>k</i>	5	0.104	0.133	0.160
	10	0.211	0.200	0.263

5. DISCUSSION

Our results suggest that most of the information is obtained from the GoogLeNet embeddings. In CoView, there are many edited videos with background music, which is a potential reason why multisensory CAMs as well as affect, which were trained on datasets with no background music, affect model performance.

When we select videos with social settings as a target set for the model fine-tuning, we see increases in P@5 and P@10, which are reflected in the selection of high-quality segments. During qualitative analysis, we observed that visually appealing clips with objects and/or people and positive emotional moments such as cheering or positive responses are preferred by the model.

6. CONCLUSIONS

In this work, we present an audiovisual model to detect highlights in videos. We use video summarization labels, and we find that visual features carry most of the information for highlight detection in CoView dataset. Using a subset of the proposed features, a combination of both audio and visual features, allows to obtain better performance in the selected metrics. Similar to other works that combine lexical information with audio and video [18], as future work we propose to explore the video transcription text as additional feature to improve performance on the task. We also propose to explore end-to-end model training on a larger dataset to adapt embeddings specifically for the highlight detection task.

7. REFERENCES

- [1] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool, "Creating summaries from user videos," in *European conference on computer vision*. Springer, 2014, pp. 505–520.
- [2] Britta Wrede and Elizabeth Shriberg, "Spotting "hot spots" in meetings: Human judgments and prosodic cues," in *Eighth European Conference on Speech Communication and Technology*, 2003.
- [3] Catherine Lai, Jean Carletta, Steve Renals, K Evanini, and K Zechner, "Detecting summarization hot spots in meetings using group level involvement and turn-taking features.," in *INTERSPEECH*, 2013, pp. 2723–2727.
- [4] Chong-Wah Ngo and Feng Wang, *Video Summarization*, pp. 3320–3324, Springer US, Boston, MA, 2009.
- [5] Ana Garcia del Molino and Michael Gygli, "Phd-gifs: Personalized highlight detection for automatic gif creation," in *Proceedings of the 26th ACM international conference on Multimedia*, 2018, pp. 600–608.
- [6] Ba Tu Truong and Svetha Venkatesh, "Video abstraction: A systematic review and classification," *ACM transactions on multimedia computing, communications, and applications (TOMM)*, vol. 3, no. 1, pp. 3–es, 2007.
- [7] Michael A Smith and Takeo Kanade, "Video skimming and characterization through the combination of image and language understanding techniques," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 1997, pp. 775–781.
- [8] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman, "Video summarization with long short-term memory," in *European conference on computer vision*. Springer, 2016, pp. 766–782.
- [9] Kaiyang Zhou, Yu Qiao, and Tao Xiang, "Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [10] Jiri Fajtl, Hajar Sadeghi Sokeh, Vasileios Argyriou, Dorothy Monekoso, and Paolo Remagnino, "Summarizing videos with attention," in *Asian Conference on Computer Vision*. Springer, 2018, pp. 39–54.
- [11] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes, "TVSum: Summarizing web videos using titles," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 5179–5187.
- [12] Yudong Jiang, Kaixu Cui, Bo Peng, and Changliang Xu, "Comprehensive video understanding: Video summarization with content-based video recommender design," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [13] Huan Yang, Baoyuan Wang, Stephen Lin, David Wipf, Minyi Guo, and Baining Guo, "Unsupervised extraction of video highlights via robust recurrent auto-encoders," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4633–4641.
- [14] Kongwah Wan and Changsheng Xu, "Efficient multimodal features for automatic soccer highlight generation," in *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004*. IEEE, 2004, vol. 3, pp. 973–976.
- [15] Michele Merler, Khoi-Nguyen C Mac, Dhiraj Joshi, Quoc-Bao Nguyen, Stephen Hammer, John Kent, Jinjun Xiong, Minh N Do, John R Smith, and Rogerio Schmidt Feris, "Automatic curation of sports highlights using multimodal excitement features," *IEEE Transactions on Multimedia*, vol. 21, no. 5, pp. 1147–1160, 2018.
- [16] Serhan Dagtas and Mohamed Abdel-Mottaleb, "Multimodal detection of highlights for multimedia content," *Multimedia Systems*, vol. 9, no. 6, pp. 586–593, 2004.
- [17] Theodoros Kostoulas, Guillaume Chanel, Michal Muszynski, Patrizia Lombardo, and Thierry Pun, "Dynamic time warping of multimodal signals for detecting highlights in movies," in *Proceedings of the 1st Workshop on Modeling INTERPERSONAL Synchrony And influence*, 2015, pp. 35–40.
- [18] Dave Makhervaks, William Hinthorn, Dimitrios Dimitriadis, and Andreas Stolcke, "Combining acoustics, content and interaction features to find hot spots in meetings," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8054–8058.
- [19] "CoVieW 2019: Comprehensive Video Understanding in the Wild," <http://cvlab.hanyang.ac.kr/coview2019/>, Accessed: 2020-07-01.
- [20] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [22] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba, "Places: A 10 million image database for scene recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [23] Dimitrios Kollias and Stefanos Zafeiriou, "Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcfac," *arXiv preprint arXiv:1910.04855*, 2019.
- [24] Andrew Owens and Alexei A Efros, "Audio-visual scene analysis with self-supervised multisensory features," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 631–648.
- [25] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.
- [26] G. Richard, S. Sundaram, and S. Narayanan, "An overview on perceptually motivated audio indexing and classification," *Proceedings of the IEEE*, vol. 101, no. 9, pp. 1939–1954, 2013.
- [27] T. Yao, T. Mei, and Y. Rui, "Highlight detection with pairwise deep ranking for first-person video summarization," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 982–990.
- [28] Kalervo Järvelin and Jaana Kekäläinen, "Cumulated gain-based evaluation of ir techniques," *ACM Transactions on Information Systems (TOIS)*, vol. 20, no. 4, pp. 422–446, 2002.