

SOUNDING HIGHLIGHTS: DUAL-PATHWAY AUDIO ENCODERS FOR AUDIO-VISUAL VIDEO HIGHLIGHT DETECTION

Seohyun Joo^{1*} Yoori Oh²

¹ School of Electrical Engineering and Computer Science, GIST

² Music and Audio Research Group, Seoul National University
seohyunj@gm.gist.ac.kr, yoori0203@snu.ac.kr

ABSTRACT

Audio-visual video highlight detection aims to automatically identify the most salient moments in videos by leveraging both visual and auditory cues. However, existing models often underutilize the audio modality, focusing on high-level semantic features while failing to fully leverage the rich, dynamic characteristics of sound. To address this limitation, we propose a novel framework, Dual-Pathway Audio Encoders for Video Highlight Detection (DAViHD). The dual-pathway audio encoder is composed of a semantic pathway for content understanding and a dynamic pathway that captures spectro-temporal dynamics. The semantic pathway extracts high-level information by identifying the content within the audio, such as speech, music, or specific sound events. The dynamic pathway employs a frequency-adaptive mechanism as time evolves to jointly model these dynamics, enabling it to identify transient acoustic events via salient spectral bands and rapid energy changes. We integrate the novel audio encoder into a full audio-visual framework and achieve new state-of-the-art performance on the large-scale Mr.HiSum benchmark. Our results demonstrate that a sophisticated, dual-faceted audio representation is key to advancing the field of highlight detection.

Index Terms— Video Highlight Detection, Audio-Visual Learning, Multimodal Fusion, Spectro-Temporal Dynamics

1. INTRODUCTION

With the rapid growth of digital video, automated video understanding has become an essential area of research. A key task within this field is video highlight detection, which focuses on identifying the most interesting and engaging moments from a video. This capability is fundamental for applications such as effective content summarization, retrieval, and recommendation.

Initial research in video highlight detection primarily centered on visual-only approaches like PGL-SUM [1]. Subsequent work began to integrate audio, with methods like Joint-VA [2] employing bimodal attention to model cross-modal interactions. More recently, the field has shifted towards large-scale, Transformer-based architectures. While models like VATT [3], MBT [4], and UMT [5] introduced sophisticated multimodal fusion mechanisms, they often treat audio simplistically, relying on generic high-level semantic features from pre-trained backbones like PANNs [6] while overlooking its rich spectro-temporal dynamics. This represents a significant gap, as the acoustic properties of sound itself are critical for identifying salient moments. The significance of these dynamics is particularly apparent in video highlight detection in Figure 1. Baseline

* This work was done during an internship at Music and Audio Research Group, Seoul National University.

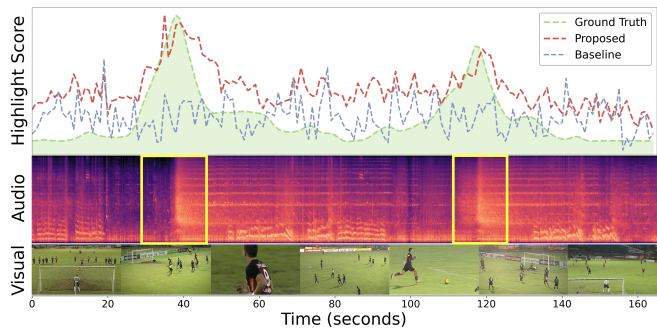


Fig. 1: Comparison of baseline and proposed framework. The **baseline model** [2] produces uniform scores by relying on global audio-visual features, failing to match the **ground-truth**. Our **proposed framework**, however, accurately captures the **ground-truth** dynamics by utilizing abrupt auditory changes, highlighted by the yellow boxes in the audio, as key features.

approaches (blue) are insensitive to the transient acoustic events that signal key moments due to their reliance on general audio-visual information. Our framework (red), however, leverages the abrupt auditory changes visualized in the spectrogram’s yellow boxes as crucial clues. By modeling these specific dynamics, it can predict highlight scores that closely mirror the green ground-truth pattern.

To address this, we draw inspiration from advances in audio representation learning. Early convolution-based methods [7, 8] explored temporal properties but were limited by the spatial equivalence assumption of 2D convolutions. A key breakthrough was Frequency-Dynamic Convolution (FDC) [9], which was proposed to address the physical inconsistencies of applying standard CNNs to spectrograms. Building on subsequent refinements of FDC that aim to capture key spectro-temporal information for transient acoustic events [10, 11], we propose a framework that utilizes two critical facets of sound: the semantic content of sound (*i.e.*, *what* is heard) and its spectro-temporal dynamics (*i.e.*, *how* the sound evolves).

To this end, we introduce a novel framework, Dual-Pathway Audio Encoders for Video Highlight Detection (DAViHD). Our approach explicitly models both the high-level semantic content and the low-level spectro-temporal dynamics of the audio signal. The main contributions of our work are as follows:

- We propose the Dual-Pathway Audio Encoder, an architecture that disentangles the audio signal into two distinct streams to model both its semantic and dynamic properties.
- By integrating our audio encoder into a unified audio-visual framework, we achieve state-of-the-art (SOTA) performance on the large-scale Mr.HiSum [12] benchmark.

The demo is available at <https://seohyj.github.io/soundhd.github.io/>.

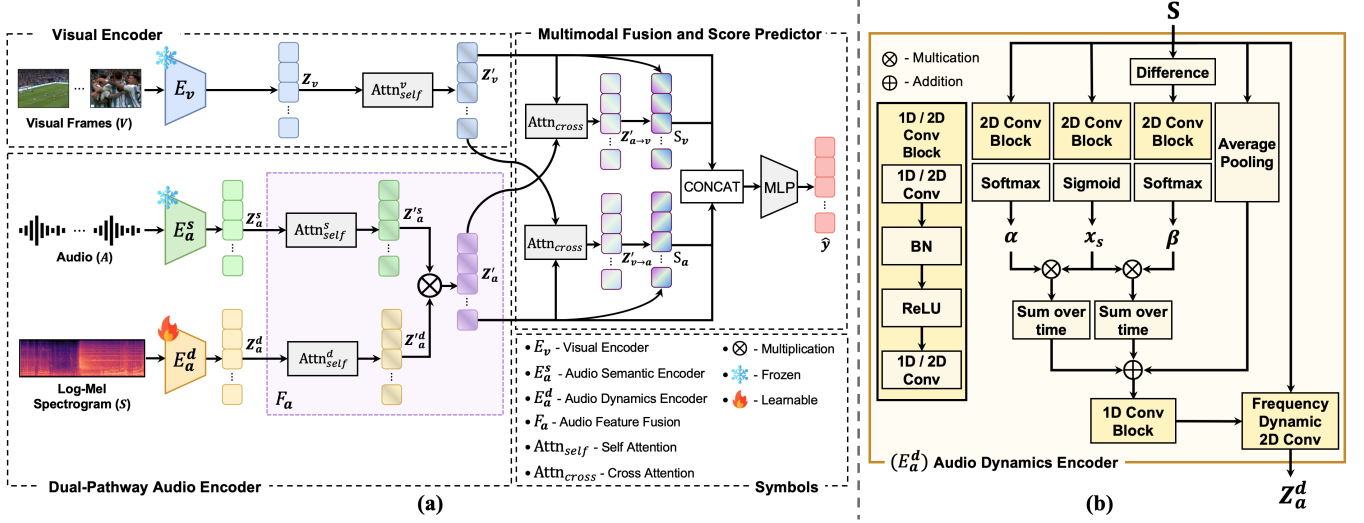


Fig. 2: (a) An overview of the DAViHD framework. The model comprises a Visual Encoder (E_v) and a Dual-Pathway Audio Encoder (E_a^s, E_a^d). Features from both audio encoders are fused via the Audio Feature Fusion module (F_a), and the fused audio feature, Z'_a , is passed to a cross-attention module and an MLP to predict the final score \hat{y} . (b) Detailed architecture of the Audio Dynamics Encoder (E_a^d). It uses a multi-branch architecture to fuse two attention maps (α, β), modulated by a saliency gate (x_s), with a global average-pooled feature. This information is then used to dynamically control a Frequency-Dynamic convolutional layer.

2. METHODS

Figure 2 illustrates the architecture of our proposed DAViHD framework. The framework processes visual and audio streams using a visual encoder and our novel dual-pathway audio encoder, respectively. The extracted multi-modal features are subsequently fused and fed into a score predictor to yield the final highlight scores.

2.1. Visual Encoder

The visual stream takes a sequence of video frames $\mathbf{V} \in \mathbb{R}^{T_f \times H \times W \times C}$ as input, where T_f , H , W , and C denote the number of frames, height, width, and channels, respectively. Initially, a pre-trained convolutional neural network, E_v , is employed as the visual backbone [13, 14] to extract a sequence of frame-level features, $\mathbf{Z}_v = E_v(\mathbf{V})$. The resulting feature sequence $\mathbf{Z}_v \in \mathbb{R}^{T_f \times D_v}$, where D_v is the visual embedding dimension, is then processed by a multi-head self-attention mechanism [15] to capture long-range contextual dependencies, yielding the final visual representation $\mathbf{Z}'_v \in \mathbb{R}^{T_f \times D_v}$.

2.2. Dual-Pathway Audio Encoder

As illustrated in Figure 2(a), the dual-pathway audio encoder extracts a comprehensive audio representation by processing the audio stream through two parallel pathways. The first pathway, an audio semantic encoder, processes the raw audio waveform in 1-second segments to model high-level semantic properties. Concurrently, the second pathway, an audio dynamics encoder, takes a log-Mel spectrogram as input to capture low-level dynamic properties. The outputs from these pathways, audio semantic embeddings $\mathbf{Z}_a^s \in \mathbb{R}^{T_f \times D_s}$ and audio dynamics features $\mathbf{Z}_a^d \in \mathbb{R}^{T_f \times D_d}$, are then fused by an audio feature fusion module. This creates a unified feature sequence $\mathbf{Z}_a \in \mathbb{R}^{T_f \times D_a}$, where D_s, D_d , and D_a are the respective dimensions of the semantic, dynamic, and fused audio embeddings.

2.2.1. Audio Semantic Encoder

The audio semantic encoder, E_a^s , is designed to interpret the high-level meaning of the audio signal, essentially, *what* is being heard. To achieve this, we employ the pre-trained PANN model [6], which

takes the raw audio waveform $\mathbf{A} \in \mathbb{R}^L$ as input, where L denotes the number of samples. The model first segments the waveform into a sequence of non-overlapping 1-second chunks. Each chunk is then processed independently by the encoder to extract a high-dimensional semantic embedding. Finally, these segment-level embeddings are concatenated chronologically to form the final sequence of semantic embeddings $\mathbf{Z}_a^s \in \mathbb{R}^{T_f \times D_s}$. We denote this entire procedure as $\mathbf{Z}_a^s = E_a^s(\mathbf{A})$.

2.2.2. Audio Dynamics Encoder

A standard convolutional block is utilized throughout the model, termed the 1D/2D Conv block. It is defined as a sequence of a Convolutional layer (1D or 2D), Batch Normalization, a ReLU activation, and a final, dimension-matching Convolutional layer. The audio dynamics encoder, E_a^d , is designed to capture low-level, time-varying patterns from a log-Mel spectrogram, $\mathbf{S} \in \mathbb{R}^{F \times T}$, where F is the number of frequency bins and T is the total number of time frames indexed by t . The encoder generates a final audio dynamics feature, $\mathbf{Z}_a^d = E_a^d(\mathbf{S})$. As shown in Figure 2(b), it processes the input through parallel branches. Two branches take the log-Mel spectrogram \mathbf{S} as input and use 2D Conv blocks to generate a temporal attention map α via a softmax function and a saliency gate x_s via a sigmoid function. A third branch captures temporal changes by processing the frame-to-frame difference, $\Delta \mathbf{S} = |\mathbf{S}_t - \mathbf{S}_{t-1}|$, which is calculated by a Difference module. $\Delta \mathbf{S}$ is then passed through the branch's own 2D Conv block and a softmax function to produce a velocity attention map, β . These maps are then used to compute time-aware features (\mathbf{f}_{TA}) and velocity-aware features (\mathbf{f}_{VA}):

$$\mathbf{f}_{TA} = \sum_t (\alpha \otimes \mathbf{x}_s), \quad \mathbf{f}_{VA} = \sum_t (\beta \otimes \mathbf{x}_s), \quad (1)$$

where \otimes denotes element-wise multiplication. These features are then summed with a global context vector from a fourth average pooling branch to form a final combined vector, $\mathbf{f}_{combined}$. This vector controls a Frequency-Dynamic 2D Conv layer [9]. The combined vector $\mathbf{f}_{combined}$ is passed through a 1D Conv Block to generate frequency-specific weights $\gamma_k(f)$, where f denotes a specific fre-

quency bin, for a set of K learnable basis kernels $\{\mathbf{W}_k\}_{k=1}^K$. These weights modulate the basis kernels, and the resulting dynamic filter is applied to the original input spectrogram \mathbf{S} . The final audio dynamics feature \mathbf{Z}_a^d is computed as the following Eq. (2), where $*$ denotes the 2D convolution.

$$\mathbf{Z}_a^d = \sum_{k=1}^K \gamma_k \otimes (\mathbf{W}_k * \mathbf{S}). \quad (2)$$

This mechanism, unlike standard 2D convolutions that treat the time and frequency axes as spatially equivalent, allows the layer to form an adaptive filter by dynamically combining the outputs of the fixed basis kernels. The feature map output from the Frequency Dynamic 2D Conv layer is transformed into the final sequence representation. The channel and frequency dimensions of the feature map are first flattened, then temporally aligned with the other modalities via adaptive average pooling to match the temporal length of T_f . Finally, a linear projection layer maps the sequence to the target embedding dimension, yielding the final audio dynamics representation \mathbf{Z}_a^d .

2.2.3. Audio Feature Fusion

The audio feature fusion module F_a in Figure 2(a) integrates the representations from the semantic and dynamics pathways. To contextualize each stream prior to fusion, the initial semantic feature, \mathbf{Z}_a^s , and dynamics feature, \mathbf{Z}_a^d , are each passed through self-attention [15]. This approach of applying self-attention before fusion, which we term Early-SA, allows the model to independently capture the temporal dependencies in each stream. As will be demonstrated in our ablation study (Section 3.3), this prior contextualization is crucial for effective fusion. This operation yields refined, context-aware representations $\mathbf{Z}_a^{s'} = \text{Attn}_{self}^s(\mathbf{Z}_a^s)$ and $\mathbf{Z}_a^{d'} = \text{Attn}_{self}^d(\mathbf{Z}_a^d)$, where Attn_{self}^s and Attn_{self}^d denote two independent self-attention blocks with their own learnable parameters. These contextualized features are then fused using element-wise multiplication. This fusion strategy acts as a gating mechanism, allowing the low-level dynamics features ($\mathbf{Z}_a^{d'}$) to modulate the high-level semantic features ($\mathbf{Z}_a^{s'}$) and amplify the salient information. The final unified audio representation, $\mathbf{Z}'_a \in \mathbb{R}^{T_f \times D_a}$, is computed as:

$$\mathbf{Z}'_a = \mathbf{Z}_a^{s'} \otimes \mathbf{Z}_a^{d'}, \quad (3)$$

where \otimes denotes element-wise multiplication. This representation is then passed directly to the multimodal fusion stage.

2.3. Multimodal Fusion and Score Prediction

The final stage of our framework fuses the context-aware visual and audio representation, \mathbf{Z}'_v , and \mathbf{Z}'_a . We employ a bidirectional cross-modal attention Attn_{cross} , following prior works [2–5, 16], to allow each modality to attend to and incorporate features from the other. First, we compute an audio-contextualized visual representation, $\mathbf{Z}'_{a \rightarrow v}$, where the visual representations act as queries (\mathbf{Q}_v) and the audio representations provide the keys (\mathbf{K}_a) and values (\mathbf{V}_a):

$$\mathbf{Z}'_{a \rightarrow v} = \text{softmax} \left(\frac{\mathbf{Q}_v \mathbf{K}_a^T}{\sqrt{d_k}} \right) \mathbf{V}_a. \quad (4)$$

Symmetrically, we compute a visual-contextualized audio representation, $\mathbf{Z}'_{v \rightarrow a}$, by swapping the roles of the audio and visual streams. To preserve original modality-specific information during this fusion, we then apply a residual connection [13] to each stream, resulting in the enhanced representations $\mathbf{S}_v = \mathbf{Z}'_v + \mathbf{Z}'_{a \rightarrow v}$ and $\mathbf{S}_a = \mathbf{Z}'_a + \mathbf{Z}'_{v \rightarrow a}$.

For the final prediction, the original self-attended features (\mathbf{Z}'_v , \mathbf{Z}'_a) and the enhanced cross-attended features (\mathbf{S}_v , \mathbf{S}_a) are all concatenated together. The resulting concatenated vector is processed by a 3-layer MLP which regresses the final frame-level highlight score, normalized between 0 and 1. The model is optimized by minimizing the Mean Squared Error (MSE) between the predicted scores \hat{y} and the ground-truth scores y :

$$\mathcal{L}_{\text{MSE}} = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{y}_t)^2. \quad (5)$$

3. EXPERIMENTS

3.1. Experimental Setups

3.1.1. Datasets

The performance is evaluated on the TVSum [17] and Mr.HiSum [12] video highlight detection benchmarks. TVSum consists of 50 web videos from 10 diverse categories. Following standard evaluation protocols [1, 18], we use 5-fold cross-validation and report the average performance. Mr.HiSum is a large-scale dataset comprising 31,892 YouTube videos with an average length of 201.9 seconds. Its highlight scores are derived from YouTube’s ‘Most replayed’ statistics, providing a robust, user-driven measure. After filtering for unavailable videos, the experiments are conducted on 30,656 videos.

3.1.2. Evaluation Metrics

We use a set of metrics to assess different aspects of its performance. To measure temporal localization accuracy, we use the F1-score. Following standard evaluation protocols, we convert the continuous prediction scores into a binary summary by selecting the top 50% of segments and measure its temporal intersection with the ground-truth summary. For ranking quality, we report mean Average Precision (mAP). We compute mAP at two thresholds, $\text{mAP}_{\rho=15\%}$ and $\text{mAP}_{\rho=50\%}$, which evaluate whether ground-truth highlight segments are ranked within the top 15% and 50% of all predictions. Finally, to evaluate the alignment with human perception, we use Spearman’s ρ and Kendall’s τ rank correlation coefficients. These metrics directly measure the monotonic agreement between the predicted score sequence and the ground-truth annotations.

3.1.3. Implementation Details

All input videos are processed into 1-second segments at a rate of 1 fps. For the visual stream, the feature dimension D_v is adapted for each benchmark to align with standard practices. On the TVSum dataset, we follow prior work [2, 20] and extract 512-dimensional visual features ($D_v = 512$) for each segment using a 3D CNN [21] with a ResNet-34 [13] backbone pre-trained on Kinetics-400 [22]. For the Mr.HiSum dataset [12], we adhere to its standard protocol, extracting 1024-dimensional visual features ($D_v = 1024$) from an Inception-v3 [14] model pre-trained on ImageNet [23], followed by PCA. For the audio stream, our dual-pathway audio encoder processes two parallel inputs. The audio semantic encoder utilizes PANNs [6] pre-trained on AudioSet [24] to extract 2048-dimensional semantic embeddings ($D_s = 2048$). The Dynamic Pathway takes a log-Mel spectrogram (16 kHz, 2048 FFT, 256 hop, 128 mel frequency bins) as input. This spectrogram is fed into the Audio Dynamics Encoder (2.2.2), which internally downsamples the feature representation using $K = 4$ learnable basis kernels, to a final resolution of 1 fps to align with the semantic features, yielding a 2048-dimensional dynamic feature sequence ($D_d = 2048$). Finally, the semantic and dynamic features are fused, resulting in a unified audio representation with a final dimension of $D_a = 2048$. Our model is trained using the Adam optimizer [25]. For the Mr.HiSum

Table 1: Main results and comparison on the Mr.HiSum and TVSum datasets. All metrics are reported as mean \pm std. over 5 runs. Higher values indicate better performance, and the best results for each metric are highlighted in bold.

Model	Mr.HiSum					TVSum				
	F1 \uparrow	mAP $_{\rho=50\%}$ \uparrow	mAP $_{\rho=15\%}$ \uparrow	ρ \uparrow	τ \uparrow	F1 \uparrow	mAP $_{\rho=50\%}$ \uparrow	mAP $_{\rho=15\%}$ \uparrow	ρ \uparrow	τ \uparrow
PGL-SUM [†] [1]	53.34 \pm 0.10	59.73 \pm 0.17	25.71 \pm 0.30	0.104 \pm 0.003	0.070 \pm 0.002	52.93 \pm 1.75	56.68 \pm 2.33	23.18 \pm 1.96	0.056 \pm 0.040	0.038 \pm 0.027
CSTA [†] [19]	54.32 \pm 0.17	61.12 \pm 0.39	28.35 \pm 0.48	0.138 \pm 0.005	0.095 \pm 0.004	57.32 \pm 1.99	62.36 \pm 2.81	27.52 \pm 5.08	0.205\pm0.056	0.141\pm0.041
Joint-VA [‡] [2]	54.71 \pm 0.04	61.82 \pm 0.11	29.09 \pm 0.22	0.152 \pm 0.001	0.104 \pm 0.001	55.03 \pm 2.20	60.94 \pm 3.19	26.66 \pm 3.40	0.142 \pm 0.046	0.097 \pm 0.031
UMT [‡] [5]	58.18 \pm 0.29	65.81 \pm 0.31	33.79 \pm 0.35	0.239 \pm 0.006	0.174 \pm 0.004	57.54 \pm 0.87	61.49 \pm 2.91	25.24 \pm 5.05	0.175 \pm 0.022	0.121 \pm 0.015
DAViHD (Ours)[‡]	59.73\pm0.41	67.27\pm0.52	36.55\pm0.51	0.299\pm0.012	0.213\pm0.009	57.67\pm1.27	63.52\pm2.58	28.94\pm3.11	0.200 \pm 0.032	0.138 \pm 0.022

[†] video only, [‡] video and audio

Table 2: Ablation study results for each modality’s contribution. V denotes the visual modality, while A_s and A_d represent the audio semantic and dynamics pathways, respectively.

V	A_s	A_d	F1 \uparrow	mAP $_{\rho=50\%}$ \uparrow	mAP $_{\rho=15\%}$ \uparrow	ρ \uparrow	τ \uparrow
✓			52.98	58.93	25.31	0.101	0.069
	✓		53.25	60.11	28.21	0.109	0.075
		✓	57.53	63.88	33.15	0.244	0.175
✓	✓		54.79	61.95	28.94	0.153	0.105
✓		✓	58.25	65.84	35.51	0.269	0.191
	✓	✓	59.09	66.12	35.62	0.282	0.203
✓	✓	✓	60.17	68.01	36.96	0.312	0.224

Table 3: Ablation study on different audio fusion strategies. We analyze the impact of the self-attention (SA) layer placement (Early vs. Late) and the feature combination method (Concat vs. Multiply).

SA Placement	Combination	F1 \uparrow	mAP $_{\rho=50\%}$ \uparrow	mAP $_{\rho=15\%}$ \uparrow	ρ \uparrow	τ \uparrow
Late	Concat	58.71	66.24	35.61	0.280	0.198
Late	Multiply	58.40	66.01	35.93	0.276	0.195
Early	Concat	59.42	67.36	36.21	0.294	0.208
Early	Multiply	60.17	68.01	36.96	0.312	0.224

dataset, we train for 200 epochs with a learning rate of 1×10^{-5} and a batch size of 16. For TVSum dataset, we train for 400 epochs with a learning rate of 5×10^{-6} and a batch size of 8. For both datasets, we use a weight decay of 1×10^{-4} and apply gradient clipping with a maximum norm of 0.5.

3.2. Main Results

Table 1 compares the performance of the proposed DAViHD model against several highlight detection methods, including PGL-SUM [1], Joint-VA [2], UMT [5] and CSTA [19]. The results indicate that our model outperforms the baseline models in most evaluation metrics, achieving SOTA performance especially on the large-scale Mr.HiSum dataset. To demonstrate the benefits of the audio modality, the DAViHD model is compared against a strong vision-only model (PGL-SUM and CSTA). Furthermore, it is benchmarked against two recent audio-visual models (Joint-VA and UMT). Joint-VA is considered a primary baseline due to its structural similarity to our approach. For UMT, its text-search functionality was disabled during evaluation to ensure a fair comparison. To ensure the reliability of the results, each experiment was conducted five times with different random initializations, and the reported scores represent the average of these runs. Notably, the proposed model outperforms even the powerful audio-visual baselines. This result highlights the effectiveness of the dual-pathway audio encoder in extracting and utilizing rich, informative features from the audio stream to identify salient moments in videos. A detailed analysis of each model component is presented in the following section.

3.3. Ablation Studies

We conduct a series of ablation studies to analyze the contribution of each component in our proposed framework. All ablation experi-

ments are performed on the large-scale, Mr.HiSum dataset.

Contribution of each Modality We evaluate the effectiveness of each input modality in an ablation study. As shown in Table 2, our ablation study reveals several key insights into the contribution of each modality. First, when considering single streams, the audio dynamics pathway (A_d) alone significantly outperforms configurations using only the visual (V) or audio semantic (A_s) streams, demonstrating that spectro-temporal dynamics are highly effective for identifying highlights. Furthermore, the combination of our two audio streams ($A_s + A_d$) yields remarkable performance. This audio-only configuration is nearly on par with the proposed model ($V + A_s + A_d$) and substantially surpasses the conventional bimodal combination of $V + A_s$. These findings strongly indicate that a sophisticated audio representation that integrates both semantic and dynamic features can provide cues for highlight detection.

Audio Fusion Strategy We examine the key design choices within our Audio Feature Fusion module (described in Section 2.2.3) to verify the effectiveness of our proposed architecture. Our model employs an Early-SA configuration, where self-attention is applied to each auditory stream before fusion, combined with element-wise Multiplication for the fusion operation itself (Eq. 3). We compare this design against three alternatives: Early-SA with Concatenation, and two Late-SA configurations with Concatenation and Multiplication. In the Late-SA configuration, self-attention is applied single time, after the initial semantic (Z_a^s) and dynamic (Z_a^d) features are fused. As shown in Table 3, the results confirm that our proposed configuration (Early-SA with Multiplication) achieves the best performance. The most critical factor is the SA placement, with Early-SA consistently outperforming Late-SA. This finding suggests that modeling temporal context within each audio stream individually is crucial before they are fused. Furthermore, the superiority of Multiplication over Concatenation within the effective Early-SA setting can be attributed to its function as a gating mechanism. A key factor in our model’s success is the synergy between pathways, wherein dynamic features effectively modulate semantic features to produce a more discriminative representation. This mechanism allows the model to amplify semantic content that is temporally correlated with salient dynamic events, leading to a more robust highlight detection.

4. CONCLUSION

In this work, we introduced a novel Dual-Pathway Audio Encoders for Audio-Visual Video Highlight Detection. Our architecture is designed to explicitly model both the high-level semantic content and the low-level spectro-temporal dynamics of the audio signal, addressing a key limitation in prior works where audio is often treated superficially. Our experiments demonstrate the effectiveness of our approach, achieving new state-of-the-art performance on the large-scale Mr.HiSum benchmark. Notably, our audio-only configuration surpassed conventional audio-visual models, underscoring the importance of a sophisticated audio representation.

Acknowledgements. This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT)[No. RS-2022-II220320, 20220-00320, 50%] and [No. RS2022-II220641, 50%].

5. REFERENCES

- [1] Evlampios Apostolidis, Georgios Balaouras, Vasileios Mezaris, and Ioannis Patras, “Combining global and local attention with positional encoding for video summarization,” in *2021 IEEE international symposium on multimedia (ISM)*. IEEE, 2021, pp. 226–234.
- [2] Taivanbat Badamdorj, Mrigank Rochan, Yang Wang, and Li Cheng, “Joint visual and audio learning for video highlight detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 8127–8137.
- [3] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong, “Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text,” *Advances in neural information processing systems*, vol. 34, pp. 24206–24221, 2021.
- [4] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun, “Attention bottlenecks for multimodal fusion,” *Advances in neural information processing systems*, vol. 34, pp. 14200–14213, 2021.
- [5] Ye Liu, Siyuan Li, Yang Wu, Chang-Wen Chen, Ying Shan, and Xiaohu Qie, “Umt: Unified multi-modal transformers for joint video moment retrieval and highlight detection,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 3042–3051.
- [6] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley, “Panns: Large-scale pre-trained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [7] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer, “Audiovisual slowfast networks for video recognition,” *arXiv preprint arXiv:2001.08740*, 2020.
- [8] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen, “Slow-fast auditory streams for audio recognition,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 855–859.
- [9] Hyeonuk Nam, Seong-Hu Kim, Byeong-Yun Ko, and Yong-Hwa Park, “Frequency dynamic convolution: Frequency-adaptive pattern recognition for sound event detection,” in *Proc. Interspeech 2022*, 2022, pp. 2763–2767.
- [10] Hyeonuk Nam, Seong-Hu Kim, Deokki Min, Junhyeok Lee, and Yong-Hwa Park, “Diversifying and expanding frequency-adaptive convolution kernels for sound event detection,” in *Proc. Interspeech 2024*, 2024, pp. 97–101.
- [11] Hyeonuk Nam and Yong-Hwa Park, “Temporal attention pooling for frequency dynamic convolution in sound event detection,” *arXiv preprint arXiv:2504.12670*, 2025.
- [12] Jinhwan Sul, Jihoon Han, and Joonseok Lee, “Mr. hisum: A large-scale dataset for video highlight detection and summarization,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 40542–40555, 2023.
- [13] 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.
- [14] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [16] Yan-Bo Lin, Jie Lei, Mohit Bansal, and Gedas Bertasius, “Eclipse: Efficient long-range video retrieval using sight and sound,” in *European Conference on Computer Vision*. Springer, 2022, pp. 413–430.
- [17] 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.
- [18] Min Jung Lee, Dayoung Gong, and Minsu Cho, “Video summarization with large language models,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 18981–18991.
- [19] Jaewon Son, Jaehun Park, and Kwangsu Kim, “Csta: Cnn-based spatiotemporal attention for video summarization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 18847–18856.
- [20] Lezi Wang, Dong Liu, Rohit Puri, and Dimitris N Metaxas, “Learning trailer moments in full-length movies with co-contrastive attention,” in *European Conference on Computer Vision*. Springer, 2020, pp. 300–316.
- [21] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh, “Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet?,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 6546–6555.
- [22] Joao Carreira and Andrew Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” in *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6299–6308.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [24] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [25] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.