

# Self-Attention Based Generative Adversarial Networks For Unsupervised Video Summarization

Maria Nektaria Minaidi

*School of ECE*

*National Technical University of Athens* *National Technical University of Athens* *National Technical University of Athens*

Athens, Greece

minaidimaria@gmail.com

Charilaos Papaioannou

*School of ECE*

Athens, Greece

cpapaioan@mail.ntua.gr

Alexandros Potamianos

*School of ECE*

Athens, Greece

potam@central.ntua.gr

**Abstract**—In this paper, we study the problem of producing a comprehensive video summary following an unsupervised approach that relies on adversarial learning. We build on a popular method where a Generative Adversarial Network (GAN) is trained to create representative summaries, indistinguishable from the originals. The introduction of the attention mechanism into the architecture for the selection, encoding and decoding of video frames, shows the efficacy of self-attention and transformer in modeling temporal relationships for video summarization. We propose the SUM-GAN-AED model that uses a self-attention mechanism for frame selection, combined with LSTMs for encoding and decoding. We evaluate the performance of the SUM-GAN-AED model on the SumMe, TVSum and COGNIMUSE datasets. Experimental results indicate that using a self-attention mechanism as the frame selection mechanism outperforms the state-of-the-art on SumMe and leads to comparable to state-of-the-art performance on TVSum and COGNIMUSE.

**Index Terms**—Unsupervised Video Summarization, Generative Adversarial Networks, key-frame extraction, Long Short-Term Memory, Deep Neural Networks

## I. INTRODUCTION

The amount of video data produced on a daily basis is growing at an exponential rate. Given this growth, users increasingly require assistance for selecting, browsing and consuming such extensive collections of videos. Video summarization aims to provide a short visual summary of an original, full-length video, that encapsulates the flow of the story and the most important segments of the video. The goal is for the produced summary to retain only the significant parts and contain as little unnecessary content as possible [1].

Several methods have been proposed to tackle video summarization using information extracted from the audio, video and text modalities.

Early approaches rely on the statistical processing of low-level video, audio and text features for assessing frame similarity or performing clustering-based key-frame selection [2], while the detection of the salient parts of the video is achieved using motion descriptors, color histograms and eigen-features.

Given the recent growth of neural network architectures, many deep learning based video summarization frameworks have been proposed over the last years. Deep-learning based video summarization algorithms typically represent visual content as feature vector encodings of video frames extracted from Convolutional Neural Networks (CNNs) [1]. One of the

challenges in video summarization is learning the complex temporal dependencies among the video frames. Early temporal modeling approaches used Long Short-Term Memory (LSTM) units [3], or in general, sequence-to-sequence models such as Recurrent Neural Networks (RNNs) [4], [5]. The introduction of transformers allowed for parallel computation, as well as, better modeling of the long-range temporal dependencies among the video frames [6]. Generative Adversarial Networks (GANs) [7] have also been used in video summarization algorithms. In [8], an adversarial framework is proposed consisting of a summarizer and a discriminator, both of which were based on LSTMs. GAN-based video summarization algorithms have been shown to produce state-of-the-art results [1]. Recently attention mechanisms have appeared to be effective for identifying the important parts of videos [1], [9]–[11].

In this work, we tackle video summarization as a key-segment selection problem. We adopt a GAN-based video summarization approach and build upon the SUM-GAN model proposed in [8].

Motivated by the limited memorization, long-range and temporal modeling capacity of the LSTM, as well as the success of the multi-head attention and transformer architectures in overcoming these drawbacks [6], [12], [13], we extend SUM-GAN by integrating attention mechanisms in several parts of the architecture. We perform an ablation study to identify the importance of better temporal modeling in the frame selection, encoder and decoder of SUM-GAN. We propose the SUM-GAN-AED model that relies on a pure attentional mechanism as the frame selector, while it retains the LSTM module as the encoder and decoder.

Our contributions include: 1) investigation of the efficacy of the integration of transformers into different parts of the SUM-GAN architecture, namely using a transformer for the frame selector (SUM-GAN-SAT), the encoder (SUM-GAN-STD, SUM-GAN-STSED) and the encoder-decoder (SUM-GAN-SAT, SUM-GAN-ST), 2) based on the above results, we propose the use of a self-attention mechanism for the frame selection, while retaining the LSTM architecture for the encoder and the decoder (SUM-GAN-AED), 3) the evaluation of the proposed models on the SumMe, TVSum and COGNIMUSE datasets, achieving state-of-the-art methods performance on

SumMe and competitive performance on TVSum and COGN-IMUSE. Our experiments indicate that the integration of self-attention at the frame selection stage of the architecture is more effective than at the encoder and decoder stage. Our code implementation can be found at our GitHub Repository<sup>1</sup>.

The rest of this paper is outlined as follows: Sections II-A and II-B present the baseline model and the model we propose, respectively. In Section II-C we outline the different variants of our proposed model, depending on the position of the added Transformer modules. The experiments are detailed in section III and Section IV concludes the paper.

## II. PROPOSED METHOD

### A. Baseline Model

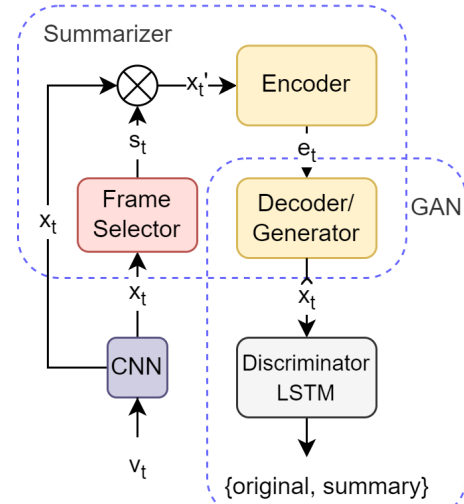
The basic generative adversarial framework [8], is illustrated in Fig. 1a. The frame selector, encoder and decoder jointly constitute the Summarizer, while the decoder (generator) along with the discriminator constitute the Generative Adversarial Network. The encoder and the decoder also form a Variational Auto-Encoder (VAE), which aids training by producing an underlying representation of the video and introducing an additional frame scores vector [14]. The approach suggests a keyframe selection mechanism, that minimizes the distance between the features of the original videos and the videos that result as a reconstruction from the predicted summaries. The summarizer and the discriminator are trained adversarially in an unsupervised manner, until the discriminator is not able to distinguish between the reconstructed videos from summaries and the original videos.

In more detail, the frame selector is a bi-directional LSTM and the encoder and decoder are LSTMs. Let  $\mathbf{X} \in \mathbb{R}^{M \times N}$  be the frame features of the input video, derived from the pretrained CNN, where  $M$  is the number of frame features and  $N$  the number of frames.  $\mathbf{x}_t \in \mathbb{R}^M$ ,  $t \in [1, N]$ , is the feature vector that describes the  $t^{\text{th}}$  frame. At each instance  $t$ , the selector is fed with  $\mathbf{x}_t$  and outputs a normalized importance scores vector  $\mathbf{s} \in \mathbb{R}^N$ , where  $\mathbf{s}_t \in [0, 1]$ . The weighted frame features,  $\mathbf{x}_t \mathbf{s}_t$ , correspond to the summary and are fed into the encoder which results in a hidden state vector,  $\mathbf{e} \in \mathbb{R}^H$ . The decoder takes  $\mathbf{e}$  as input and reconstructs a sequence of features representing the input video  $\hat{\mathbf{x}} \in \mathbb{R}^N$ . Finally,  $\hat{\mathbf{x}}$  is forwarded to the LSTM based discriminator, which aims to classify it as 'original' or 'summary'.

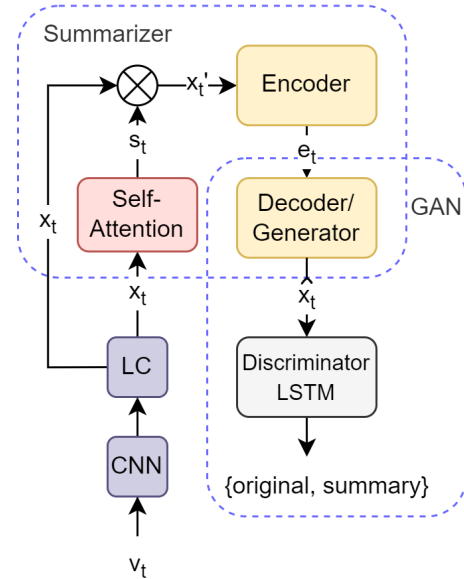
### B. The Proposed SUM-GAN-AED Model

Our approach introduces an attention-based frame selector in the GAN-based architecture. The model is inspired by [10] and [5]. The self-attention layer ranges over the entire video sequence and efficiently captures long-term temporal dependencies, as opposed to using an LSTM that fails to capture such dependencies effectively [12]. This leads to faster computation and a more representative frame selection scores vector [12].

<sup>1</sup><https://github.com/Aria-Minaidi/Self-Attention-based-GANs-for-Video-Summarization>



(a) SUM-GAN



(b) SUM-GAN-AED

Figure 1: The SUM-GAN and SUM-GAN-AED architectures.

The architecture of the proposed SUM-GAN-AED model is depicted in Fig. 1b. Following the linear compression layer, the compressed frame features sequence is forwarded to the self-attention module. The computed frame scores multiplied with the frame features are then fed into the Variational Auto-Encoder and afterwards to the discriminator. The encoder, decoder and discriminator are LSTMs as in the baseline model.

The parameters of each module of the architecture are learned during training through the loss functions that mirrors [8] and are as follows:

- The Reconstruction Loss:

$$\mathcal{L}_{reconst} = \mathbb{E}[-\log p(\phi(\hat{\mathbf{x}})|\mathbf{e})] \quad (1)$$

where  $\phi(\hat{\mathbf{x}})$  is the output of the last hidden layer of the discriminator and  $p(\phi(\hat{\mathbf{x}})|\mathbf{e}) \propto \exp(-\|\phi(\mathbf{x}) - \phi(\hat{\mathbf{x}})\|^2)$ .

- The Prior Loss:

$$\mathcal{L}_{prior} = \mathcal{D}_{KL}(q(\mathbf{e}|\mathbf{x}) \parallel \mathcal{N}(0, 1)) \quad (2)$$

where  $\mathcal{D}_{KL}$  is the Kullback–Leibler divergence and  $q(\mathbf{e}|\mathbf{x})$  is the probability of observing  $\mathbf{e}$  given  $\mathbf{x}$ .

- The Sparsity Loss:

$$\mathcal{L}_{sparsity} = \left\| \frac{1}{N} \sum_{t=1}^N s_t - \sigma \right\|_2 \quad (3)$$

where  $N$  is the number of video frames and  $\sigma \in [0, 1]$  is the summary rate that refers to the percentage of frames that constitute the summary.

- The GAN Loss:

$$\mathcal{L}_{GAN} = \log(LSTM(\mathbf{x})) + \log(1 - LSTM(\hat{\mathbf{x}})) + \log(1 - LSTM(\hat{\mathbf{x}}_p)) \quad (4)$$

where  $LSTM()$  is the softmax output of the classifier LSTM and denotes probability scores representing the discriminator’s confidence.  $\hat{\mathbf{x}}_p$  is reconstructed from the subset of video frames randomly selected by sampling from the prior distribution that was produced by the VAE.

### C. Model Variants with Self-Attention

Next, we further investigate the effectiveness of adding self-attention layers in the SUM-GAN model, as it pertains to capturing long-term temporal dependencies in video. Specifically, we experimented with replacing the LSTMs in the SUM-GAN-AED architecture with transformers at various places. The following changes are proposed: replacing the LSTM in the encoder with various flavors of a transformer (SUM-GAN-STD, SUM-GAN-STSED) and replacing the LSTMs in the encoder and the decoder with a transformer (SUM-GAN-SAT, SUM-GAN-ST). For the SUM-GAN-STD, SUM-GAN-STSED and SUM-GAN-ST models we also use a bi-directional LSTM instead of a self-attention module as the frame selector. The evaluation of the various LSTM, self-attention and transformer variations serve as an ablation study that highlights which modules in the architecture benefit most from better temporal dependency modeling. The proposed models are detailed next.

#### 1) SUM-GAN-STD

In SUM-GAN-STD we swap the encoder LSTM of [8] for a transformer [15]. The frame selector is a bi-directional LSTM and the decoder and discriminator are LSTMs as in [8]. The use of a transformer is inspired by its ability to mitigate information loss and permit parallel computations [6], [13]. Following the linear compression layer and the frame selector, the weighted feature vector is forwarded to the transformer, whose output is fed to the decoder and then the discriminator. The training follows the baseline [8].

#### 2) SUM-GAN-ST

In SUM-GAN-ST we replace the Variational Auto-Encoder, i.e. the encoder and the decoder, with a transformer, in order to enhance the video reconstruction by integrating the positional information of the frames [12]. The sequence-to-sequence architecture is suitable for video summarization. The weighted frame features enter the transformer and the reconstructed frame sequence that corresponds to the input video is fed into the discriminator, in order to be classified as ‘original’ or ‘summary’. Since we remove the VAE, during training we do not utilize the Prior Loss; the transformer is trained as part of the summarizer and the GAN [10].

#### 3) SUM-GAN-STSED

In SUM-GAN-STSED we replace the encoder of the architecture with a *Transformer Sequence Encoder* (TSE), the part of the transformer module that constitutes the encoder [15]. The intuition remains, a more effective representation of the long-range temporal video dependencies and as a result of the hidden state vector  $\mathbf{e}$ , which leads to a better video reconstruction. For its implementation, we use the slp framework<sup>2</sup>. The TSE is a sequence-to-vector architecture, which uses positional encodings to insert relative position information of the sequence tokens and keep track of the ordering of the frames. It takes the output of the frame selector multiplied by the features vector, forwards it to a linear layer, computes the positional embeddings and adds them to the tensor. The resulting vector is forwarded to the encoder part of the TSE and the output is fed to the decoder LSTM of the architecture [12]. The frame selector is an LSTM, and the training follows the description in Section II-A.

#### 4) SUM-GAN-SAT

Finally, motivated by the effectiveness of the attention based modules, we build SUM-GAN-SAT, in which we swap the LSTM frame selector with a self-attention module and the Variational Auto-Encoder with a transformer. This model combines the SUM-GAN-AED and SUM-GAN-ST architectures and their advantages in integrating attention-based modules in the GAN-based architecture. Here, we incorporate the more thorough frame selection that self-attention provides, with the integration of the positional encodings information during the video reconstruction phase [15]. We do not use Prior Loss during training, which follows the training of the preceding models.

## III. EXPERIMENTS

We evaluate the proposed model against the state-of-the-art on two benchmark datasets SumMe [16] and TVSum [17], which contain 25 and 50 videos respectively, with ground-truth annotations. For the ablation study, we add a third database COGNIMUSE [18] that consists of half-hour segments from seven Hollywood movies.

<sup>2</sup><https://github.com/georgepar/slp>

### A. Evaluation Metrics

For a fair comparison with the state-of-the-art, we report key-shot-based F-score on TVSum and SumMe datasets, as in [8]. Specifically the similarity between the automatically generated and the ground-truth summary is computed as the harmonic mean of precision and recall measured on the temporal overlap between the summaries. For the performance evaluation on COGNIMUSE, we follow the approach presented in [19]. COGNIMUSE does not include ground-truth frame scores, hence the summarization task is approached as a two-class classification problem, where multiple thresholds are applied to the estimated frame-wise importance scores. Results are obtained at various compression rates, producing summaries of various lengths. The Area Under the Curve (AUC) metric is computed using the Receiver Operating Characteristic (ROC) curve, generated at different thresholds/compression rates.

### B. Implementation Details

For all experiments, we follow the standard 5-fold cross validation approach, 80% of the videos are used for training and 20% of the videos are used for testing, and we report the average F-score over the 5 runs. We train our models over 50 epochs. We downsample the videos to 2 FPS and we use the output of pool5 layer of GoogleNet trained on ImageNet to represent the visual content of the frames. The linear compression layer reduces the size of the feature vectors from 1024 to 500. Each component of the architecture comprises a 2-layer LSTM with 500 hidden units in each layer, while the frame selector LSTM is bi-directional. Training is based on the Adam optimizer and the learning rate for all components but the discriminator is  $10^{-4}$ ; for the latter it equals to  $10^{-5}$ . The regularization factor for the sparsity loss is  $\sigma = 0.3$ . We implement our method in PyTorch.

### C. Results and Comparison

Table I depicts the F-score results of the recent state-of-the-art unsupervised video summarization methods, in comparison with our model SUM-GAN-AED. The proposed framework surpassed the state-of-the-art methods on SumMe, by a significant margin and on TVSum its performance is comparable to the current state-of-the-art. The efficacy of SUM-GAN-AED is compared to that of the experimental models, developed as part of this study, that utilize transformer architectures in Table II. Both tables showcase that the introduction of the attention mechanism, either with self-attention or a transformer, leads to an improvement of the performance on SumMe and TVSum, as well as COGNIMUSE.

Between all experimental models, SUM-GAN-AED yields the best overall results, while taking in consideration all the datasets, which also shows that the performance peak is not dataset focused. The addition of the self-attention mechanism at the frame selection stage in the transformer-based architecture SUM-GAN-ST, leads to improved performance as evidenced by SUM-GAN-SAT. This result underlines the *importance of the self-attention module in terms of overall*

Table I: Comparative performance evaluation of our SUM-GAN-AED model, with state-of-the-art unsupervised key-frame extraction approaches, on SumMe and TVSum (F-score(%)).

Model	SumMe	TVSum
SUM-GAN [8]	38.7	50.8
ACGAN [11]	46.0	58.5
Cycle-SUM [20]	46.8	57.6
SUM-GAN-sl [21]	46.8	<b>65.3</b>
SUM-GAN-AAE [10]	48.9	58.3
Proposed-B [22]	58.8	63.5
SUM-GAN-AED (Ours)	<b>64.85</b>	<b>63.18</b>

Table II: Comparative performance evaluation of our proposed models, on SumMe (F-score(%)), TVSum (F-score(%)) and COGNIMUSE (AUC).

Model	SumMe	TVSum	COGN
SUM-GAN-STD	54.15	<b>63.82</b>	51.38
SUM-GAN-ST	56.00	60.53	50.73
SUM-GAN-STSED	61.30	62.73	52.8
SUM-GAN-SAT	61.38	62.41	49.81
SUM-GAN-AED	<b>64.85</b>	<b>63.18</b>	<b>55.49</b>

*performance*. As it can be seen in Tables I and II, besides SUM-GAN-AED, our transformer-based models outperform other methods in the literature.

## IV. CONCLUSIONS

We propose a novel framework for unsupervised video summarization based on a Generative Adversarial Network. Building on the SUM-GAN model, we utilize the attention mechanism, with the introduction of self-attention and transformer modules into our framework, in order to capture long-range dependencies, adapt to sequence lengths not encountered in training and embed the positional information of the video frames. Our experiments show that the use of self-attention for frame selection, followed by LSTMs for encoding and decoding, leads to the best overall results for the video summarization task. Furthermore, the experimental transformer-based architectures perform consistently well on the evaluation datasets and outperform many state-of-the-art unsupervised approaches. Quantitative evaluation on two public benchmark datasets (TVSum, SumMe) and further evaluation on the COGNIMUSE database, proves the efficacy of our approach.

## ACKNOWLEDGMENTS

We would like to thank Georgios Paraskevopoulos and Dimitrios Sotiriou for their valuable conversations and insightful contributions to this research. Their opinions and suggestions have greatly enhanced the quality of this work, and we are thankful for their feedback throughout the process.

## REFERENCES

- [1] E. Apostolidis, E. Adamantidou, A. I. Metsai, V. Mezaris, and I. Patras, "Video summarization using deep neural networks: A survey," CoRR, vol. abs/2101.06072, 2021
- [2] R. M. Jiang, A. H. Sadka, and D. Crookes, Advances in Video Summarization and Skimming, Springer Berlin Heidelberg, 2009.

- [3] K. Zhang, W. L. Chao, F. Sha, and K. Grauman, "Video summarization with long short-term memory," *CoRR*, vol. abs/1605.08110, 2016.
- [4] M. Sanabria, F. Precioso, and T. Menguy, "Hierarchical multimodal attention for deep video summarization," in *2020 25th International Conference on Pattern Recognition (ICPR)*, 2021.
- [5] Z. Ji, K. X., Y. Pang, and X. Li, "Video summarization with attention-based encoder-decoder networks," *CoRR*, vol. abs/1708.09545, 2017.
- [6] B. Zhao, M. Gong, and X. Li, "Hierarchical multimodal transformer to summarize videos," 2021.
- [7] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," 2014.
- [8] B. Mahasseni, M. Lam, and S. Todorovic, "Unsupervised video summarization with adversarial lstm networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] Yunjae Jung, Donghyeon Cho, Dahun Kim, Sanghyun Woo, and In So Kweon, "Discriminative feature learning for unsupervised video summarization," *CoRR*, vol. abs/1811.09791, 2018.
- [10] E. Apostolidis, E. Adamantidou, A. I. Metsai, V. Mezaris, and I. Patras, "Unsupervised video summarization via attention-driven adversarial learning," in *Proceedings of the International Conference on Multimedia Modeling (MMM)*. 2020, Springer
- [11] X. He, Y. Hua, T. Song, Z. Zhang, Z. Xue, R. Ma, N. Robertson, and H. Guan, "Unsupervised video summarization with attentive conditional generative adversarial networks," in *Proceedings of the 27th ACM International Conference on Multimedia*, New York, NY, USA, 2019, MM '19, p. 2296–2304, Association for Computing Machinery.
- [12] Manjot Bilkhu, Siyang Wang, and Tushar Dobhal, "Attention is all you need for videos: Self-attention based video summarization using universal transformers," *CoRR*, vol. abs/1906.02792, 2019.
- [13] Jeiyoon Park, Kiho Kwoun, Chanhee Lee, and Heuiseok Lim, "Multimodal frame-scoring transformer for video summarization," 2022.
- [14] D. P Kingma and M. Welling, "Auto-encoding variational bayes," 2013.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2017.
- [16] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool, "Creating summaries from user videos," in *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, Eds., Cham, 2014, pp. 505–520, Springer International Publishing.
- [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 (CVPR)*, June 2015.
- [18] Zlatintsi A., Koutras P., and Evangelopoulos G. et al., "Cognimuse: a multimodal video database annotated with saliency, events, semantics and emotion with application to summarization.," *Journal on Image and Video Processing*, vol. vol. 54 (2017).
- [19] P. Koutras and P. Maragos, "Susinet: See, understand and summarize it," *CoRR*, vol. abs/1812.00722, 2018.
- [20] Li Yuan, Francis E. H. Tay, Ping Li, Li Zhou, and Jiashi Feng, "Cycle-sum: Cycle-consistent adversarial LSTM networks for unsupervised video summarization," *CoRR*, vol. abs/1904.08265, 2019.
- [21] E. Apostolidis, A. I. Metsai, E. Adamantidou, V. Mezaris, and I. Patras, "A stepwise, label-based approach for improving the adversarial training in unsupervised video summarization," in *Proceedings of the 1st International Workshop on AI for Smart TV Content Production, Access and Delivery*, 2019.
- [22] Michail Kaseris, Ioannis Mademlis, and Ioannis Pitas, "Exploiting caption diversity for unsupervised video summarization," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 1650–1654.