

Video Summarization by Learning from Unpaired Data

Mrigank Rochan and Yang Wang
 University of Manitoba

{mrochan, ywang}@cs.umanitoba.ca

Abstract

We consider the problem of video summarization. Given an input raw video, the goal is to select a small subset of key frames from the input video to create a shorter summary video that best describes the content of the original video. Most of the current state-of-the-art video summarization approaches use supervised learning and require labeled training data. Each training instance consists of a raw input video and its ground truth summary video curated by human annotators. However, it is very expensive and difficult to create such labeled training examples. To address this limitation, we propose a novel formulation to learn video summarization from unpaired data. We present an approach that learns to generate optimal video summaries using a set of raw videos (V) and a set of summary videos (S), where there exists no correspondence between V and S . We argue that this type of data is much easier to collect. Our model aims to learn a mapping function $F: V \rightarrow S$ such that the distribution of resultant summary videos from $F(V)$ is similar to the distribution of S with the help of an adversarial objective. In addition, we enforce a diversity constraint on $F(V)$ to ensure that the generated video summaries are visually diverse. Experimental results on two benchmark datasets indicate that our proposed approach significantly outperforms other alternative methods.

1. Introduction

In recent years, there has been a phenomenal surge in videos uploaded online everyday. With this remarkable growth, it is becoming difficult for users to watch or browse these online videos efficiently. In order to make this enormous amount of video data easily browsable and accessible, we need automatic video summarization tools. The goal of video summarization is to produce a short summary video that conveys the important and relevant content of a given longer video. Video summarization can be an indispensable tool that has potential applications in a wide range of domains such as video database management, consumer video analysis and surveillance [35].

Video summarization is often formulated as a subset selection problem. In general, there are two types of subset selection in video summarization: (i) key frames selection, where the goal is to identify a set of isolated frames [8, 17, 20, 24, 25, 31, 39]; and (ii) key shots selection, where the aim is to identify a set of temporally continuous interval-based segments or subshots [16, 22, 26, 27]. In this paper, we treat video summarization as a key frame selection problem. A good summary video should contain video frames that satisfy certain properties. For example, the selected frames should capture the key content of the video [8, 13, 27]. In addition, the selected frames should be visually diverse [8, 24, 39].

Both supervised and unsupervised learning approaches have been proposed for video summarization. Most unsupervised methods [13, 14, 17, 22, 24, 27, 29, 33, 47, 31, 40] use hand-crafted heuristics to select frames in a video. The limitation of such approaches is that it is difficult to come up with the heuristics that are sufficient for generating good summary videos. In contrast, supervised methods [8, 10, 11, 38, 39, 43] learn from training data with user-generated summary videos. Each instance of the training data consists of a pair of videos – a raw input video and its corresponding ground truth summary video created by humans. From such training data, these supervised methods learn to map from a raw input video to a summary video by mimicking how humans create summary videos. Supervised methods can implicitly capture cues used by humans that are difficult to model via hand-crafted heuristics, so they tend to outperform unsupervised methods.

A major limitation of supervised video summarization is that it relies on labeled training data. Common datasets in the community are usually collected by asking human annotators to watch the input video and select the key frames or key shots. This annotation process is very expensive and time-consuming. As a result, we only have very few benchmark datasets available for video summarization in the computer vision literature. Moreover, each dataset usually only contains a small number of annotated data (see Table 1).

To address the flaws of supervised learning, we propose a new formulation of learning video summarization from

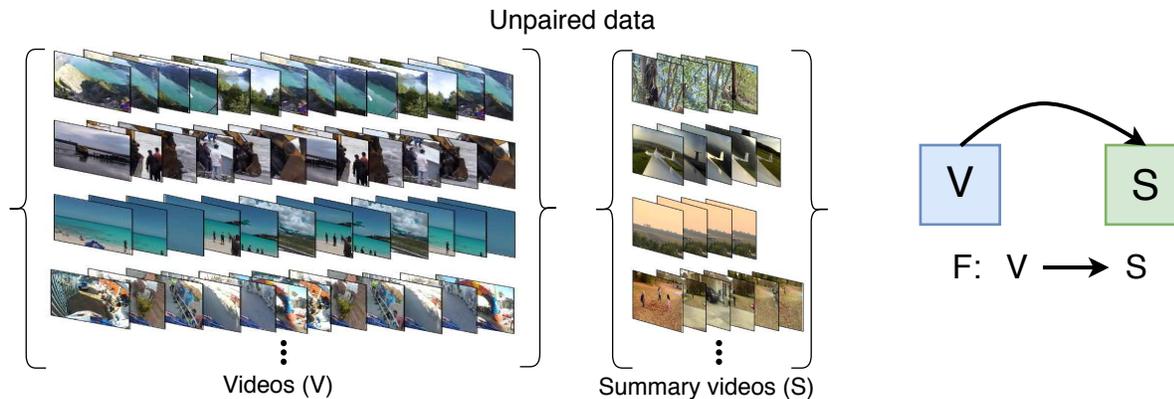


Figure 1. Learning video summarization from unpaired data. Given a set of raw videos $\{v_i\}_{i=1}^M$ ($v \in V$) and real summary videos $\{s_j\}_{j=1}^N$ ($s \in S$) such that there exists no matching/correspondence between the instances in V and S , our aim is to learn a mapping function $F: V \rightarrow S$ (right) linking two different domains V and S . The data are *unpaired* because the summary set S does not include ground truth summary videos for raw videos in V , and vice versa.

unpaired data. Our key insight is that it is much easier to collect unpaired video sets. First of all, raw videos are easily accessible as they are abundantly available on the Internet. At the same time, good summary videos are also readily available in large quantities. For example, there are lots of sports highlights, movie trailers, and other professionally edited summary videos available online. These videos can be treated as ground truth summary videos. The challenge is that these professionally curated summary videos usually do not come with their corresponding raw input videos. In this paper, we propose to solve video summarization by learning from such unpaired data (Fig. 1 (left)). We assume that our training data consist of two sets of videos: one set of raw videos (V) and another set of human created summary videos (S). However, there exists no correspondence between the videos in these two sets, i.e. the training data are *unpaired*. In other words, for a raw video in V , we may not have its corresponding ground truth summary video in S , and vice versa.

We propose a novel approach to learn video summarization from unpaired training data. Our method learns a mapping function (called the *key frame selector*) $F: V \rightarrow S$ (Fig. 1 (right)) to map a raw video $v \in V$ to a summary video $F(v)$. It also trains a *summary discriminator* that tries to differentiate between a generated summary video $F(v)$ and a real summary video $s \in S$. Using an adversarial loss [9], we learn to make the distribution of generated summary videos $F(v)$ to be indistinguishable from the distribution of real summary videos in S . As a result, the mapping function F will learn to generate a realistic summary video for a given input video. We also add more structure to our learning by introducing a reconstruction loss and a diversity loss on the output summary video $F(v)$. By combining these two losses with the adversarial loss, our method learns to generate meaningful and visually diverse

video summaries from unpaired data.

In summary, our contributions of this paper include: (i) a new problem formulation of learning video summarization from unpaired data, which consists of a set of raw videos and a set of video summaries that do not share any correspondence; (ii) a deep learning model for video summarization that learns from unpaired data via an adversarial process; (iii) an extensive empirical study on benchmark datasets to demonstrate the effectiveness of the proposed approach; and (iv) an extension of our method that introduces the partial supervision to improve the summarization performance.

2. Related Work

With the explosive increase in the amount of online video data, there has been a growing interest in the computer vision community on developing automatic video summarization techniques. Most of the prior approaches fall in the realm of unsupervised and supervised learning.

Unsupervised methods [6, 7, 12, 14, 17, 21, 22, 23, 27, 29, 30, 33, 37, 45] typically use hand-crafted heuristics to satisfy certain properties (e.g. diversity, representativeness) in order to create the summary videos. Some summarization methods also provide weak supervision through additional cues such as web images/videos [4, 13, 14, 33] and video category information [28, 30] to improve the performance.

Supervised methods [8, 10, 11, 19, 24, 31, 32, 38, 39, 40, 41, 42, 47] learn video summarization from labeled data consisting of raw videos and their corresponding ground-truth summary videos. Supervised methods tend to outperform unsupervised ones, since they can learn useful cues from ground truth summaries that are hard to capture with hand-crafted heuristics. Although supervised methods are promising, they are limited by the fact that they require ex-

pensive labeled training data in the form of videos and their summaries (i.e., paired data). In this paper, we propose a new formulation for video summarization where the algorithm only needs unpaired videos and summaries (see Fig. 1(left)) for training. The main advantage of such unpaired data is that it is much easier to collect them.

Recent methods treat video summarization as a structured prediction task [24, 31, 39, 40, 47, 43, 44, 48]. In particular, our formulation aligns with Roohan *et al.* [31] that models the video summarization as a sequence labeling problem. Unlike contemporary methods [39, 24, 43, 44, 40] that use recurrent models, Roohan *et al.* [31] propose fully convolutional sequence model which is efficient and allows better GPU parallelization. However, the major limitation of their method is that it is fully supervised and relies on paired training data. In contrast, we aim to learn video summarization using videos and summaries that have no matching information (i.e., unpaired data).

Lastly, our notion of learning from unpaired data is partly related to recent research in image-to-image translation [2, 5, 36, 49]. These methods learn to translate an input image from one domain to an output image in another domain without any paired images from both domains during training. However, there are major technical differences between these methods and ours. They typically employ two-way generative adversarial networks (GANs) with cycle consistency losses, whereas our method is an instance of standard GANs [9] with losses designed to solve video summarization. Moreover, their formulation is limited to unpaired learning in images. To the best of our knowledge, this paper is the first work on unpaired learning in video analysis, in particular video summarization.

3. Our Approach

3.1. Formulation

We are given an unpaired dataset consisting of a set of raw videos $\{v_i\}_{i=1}^M$ and a set of real summary videos $\{s_j\}_{j=1}^N$, where $v_i \in V$ and $s_j \in S$. We define the data distribution for v and s as $v \sim p_{data}(v)$ and $s \sim p_{data}(s)$, respectively. Our model consists of two sub-networks called the *key frame selector network* (S_K) and the *summary discriminator network* (S_D). The key frame selector network is a mapping function $S_K : V \rightarrow S$ between the two domains V and S (see Fig. 1). Given an input video $v \in V$, the key frame selector network (S_K) aims to select a small subset of k key frames of this video to form a summary video $S_K(v)$. The goal of the summary discriminator network (S_D) is to differentiate between a real summary video $s \in S$ and the summary video $S_K(v)$ produced by the key frame selector network S_K . Our objective function includes an adversarial loss, a reconstruction loss and a diversity loss. We learn the two networks S_K and S_D in an

adversarial fashion. In the end, S_K learns to output an optimal summary video for a given input video. In practice, we precompute the image feature of each frame in a video. With a little abuse of terminology, we use the term “video” to also denote the sequence of frame-level feature vectors when there is no ambiguity based on the context.

3.2. Network Architecture

The *key frame selector network* (S_K) in our model takes a video with T frames as the input and produces the corresponding summary video with k key frames. We use the fully convolutional sequence network (FCSN) [31], an encoder-decoder fully convolutional network, to select key frames from the input video. FCSN encodes the temporal information among the video frames by performing convolution and pooling operations in the temporal dimension. This enables FCSN to extract representations that capture the inter-frame structures. The decoder of FCSN consists of several temporal deconvolution operations which produces a vector of prediction scores with the same length as the input video. Each score indicates the likelihood of the corresponding frame being a key frame or non-key frame. Based on these scores, we select k key frames to form the predicted summary video. In order to define the reconstruction loss used in the learning (see Sec. 3.3), we apply convolution operations on the decoded feature vectors of these k key frames to reconstruct the corresponding feature vectors in the input video. We also introduce a skip connection that retrieves the frame-level feature representation of the selected k key frames, which we merge with the reconstructed features of the k key frames. Fig. 2 (a) shows the architecture of S_K .

The *summary discriminator network* (S_D) in our model takes two kinds of input: (1) the summary videos produced by S_K for the raw videos in V ; and (2) the real summary videos in S . The goal of S_D is to distinguish between the summaries produced by S_K and the real summaries. We use the encoder of FCSN [31] to encode the temporal information within the input summary video. Next, we perform a temporal average pooling operation (Ω_t) on the encoded feature vectors to obtain a video-level feature representation. Finally, we append a fully connected layer (\mathcal{FC}), followed by a sigmoid operation (σ) to obtain a score (\mathcal{R}_s) indicating whether the input summary video is a real summary or a summary produced by S_K . Let s be an input summary video to S_D , we can express the operations in S_D by Eq. 1. The network architecture of S_D is shown in Fig. 2 (b).

$$\mathcal{R}_s = S_D(s) = \sigma(\mathcal{FC}(\Omega_t(\text{FCSN}_{enc}(s)))) \quad (1)$$

3.3. Learning

Our learning objective includes an adversarial loss [9], a reconstruction loss, and a diversity loss.

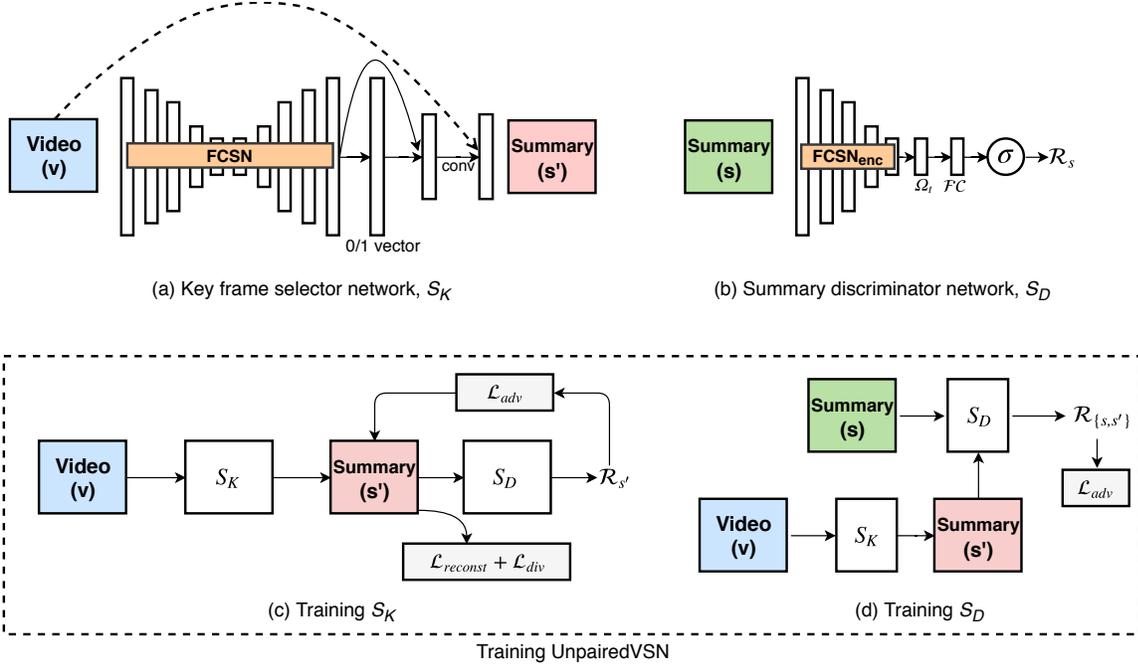


Figure 2. Overview of our proposed model. (a) Network architecture of the key frame selector network S_K . It takes a video v and produces its summary video s' (i.e., $S_K(v)$) by selecting k key frames from v . The backbone of S_K is FCSN [31]. We also introduce a skip connection from the input to retrieve the frame-level features of k key frames selected by S_K . (b) Network architecture of the summary discriminator network S_D . It differentiates between an output summary video s' and a real summary video s . S_D consists of the encoder of FCSN ($FCSN_{enc}$), followed by a temporal average pooling (Ω_t) and sigmoid (σ) operations. In (c) and (d), we show the training scheme of S_K and S_D , respectively. S_K tries to produce video summaries that are indistinguishable from real video summaries created by humans, whereas S_D tries to differentiate real summary videos from the summaries produced by S_K . As mentioned in Sec. 3.1, there is no correspondence information available to match raw videos and summary videos in the training data.

Adversarial Loss: This loss aims to match the distribution of summary videos produced by the key frame selector network S_K with the data distribution of real summary videos. We use the adversarial loss commonly used in generative adversarial networks [9]:

$$\mathcal{L}_{adv}(S_D, S_K) = \mathbb{E}_{s \sim p_{data}(s)} [\log S_D(s)] + \mathbb{E}_{v \sim p_{data}(v)} [\log(1 - S_D(S_K(v)))] \quad (2)$$

where S_K aims to produce summary videos $S_K(v)$ that are close to real summary videos in domain S , and S_D tries to differentiate between output summary videos $\{S_K(v) : v \in V\}$ and real summary videos $\{s : s \in S\}$. A minimax game occurs between S_K and S_D , where S_K pushes to minimize the objective and S_D aims to maximize it. This is equivalent to the following:

$$\min_{S_K} \max_{S_D} \mathcal{L}_{adv}(S_D, S_K) \quad (3)$$

Reconstruction Loss: We introduce a reconstruction loss to minimize the difference between the reconstructed feature representations of the k key frames in the predicted summary video $S_K(v)$ and the input frame-level representation of those k key frames in the input video v . Let Λ_k

be a set of k indices indicating which k frames in the input video are selected in the summary. In other words, if $f \in \Lambda_k$, the f -th frame in the input video is a key frame. We can define this reconstruction loss as:

$$\mathcal{L}_{reconst}(S_K(v), v) = \frac{1}{k} \sum_{t=1}^k \|S_K(v)^t - v^{f_t}\|_2^2 \quad (4)$$

where $S_K(v)^t$ and v^{f_t} are the features of the t -th frame in the output summary video $S_K(v)$ and the f_t -th frame (i.e. $f_t \in \Lambda_k$) of the input video v , respectively. The intuition behind this loss is to make the reconstructed feature vectors of the key frames in the summary video $S_K(v)$ similar to the feature vectors of those frames in the input video v .

Diversity Loss: It is desirable in video summarization that the frames in the summary video have high visual diversity [24, 39, 47]. To enforce this constraint, we apply a repelling regularizer [46] that encourages the diversity in the output summary video $S_K(v)$ for the given input video v .

This diversity loss is defined as:

$$\mathcal{L}_{div}(S_K(v)) = \frac{1}{k(k-1)} \sum_{t=1}^k \sum_{t'=1, t' \neq t}^k \frac{(S_K(v)^t)^T \cdot S_K(v)^{t'}}{\|S_K(v)^t\|_2 \|S_K(v)^{t'}\|_2} \quad (5)$$

where $S_K(v)^t$ is the frame-level reconstructed feature representation of frame t in the summary video $S_K(v)$. We aim to minimize $\mathcal{L}_{div}(S_K(v))$, so that the selected k key frames are visually diverse.

Final Loss: Our final loss function is:

$$\mathcal{L}(S_K, S_D) = \mathcal{L}_{adv}(S_D, S_K) + \mathcal{L}_{reconst}(S_K(v), v) + \beta \mathcal{L}_{div}(S_K(v)) \quad (6)$$

where β is a hyperparameter that controls the relative importance of the visual diversity. The goal of the learning is to find the optimal parameters $\Theta_{S_K}^*$ and $\Theta_{S_D}^*$ in S_K and S_D , respectively. We can express this as the following:

$$\Theta_{S_K}^*, \Theta_{S_D}^* = \arg \min_{\Theta_{S_K}, \Theta_{S_D}} \mathcal{L}(S_K, S_D) \quad (7)$$

For brevity, we use `UnpairedVSN` to denote our unpaired video summarization network that is learned by Eq. 7. In Fig. 2(c) and Fig. 2(d), we show the training scheme of S_K and S_D in our model `UnpairedVSN`.

3.4. Learning with Partial Supervision

In some cases, we may have a small amount of paired videos during training. We use V_p ($V_p \subset V$) to denote this subset of videos for which we have the ground truth summary videos. Our model can be easily extended to take advantage of this partial supervision. In this case, we apply an additional objective \mathcal{L}_{psup} on the output of FCSN in the key frame selector network S_K . Suppose a training input video $v \in V_p$ has T frames, $\delta_{t,l}$ is the score of the t -th frame to be the l -th class (key frame or non-key frame) and l_t is the ground truth binary key frame indicator. We define $\mathcal{L}_{psup}(v)$ as:

$$\mathcal{L}_{psup}(v) = -\frac{1}{T} \sum_{t=1}^T \log \left(\frac{\exp(\delta_{t,l_t})}{\sum_{l=1}^2 \exp(\delta_{t,l})} \right) \quad (8)$$

Our learning objective in this case is defined as:

$$\mathcal{L}(S_K, S_D) = \mathcal{L}_{adv} + \mathcal{L}_{reconst} + \beta \mathcal{L}_{div} + \gamma \cdot \mathbb{1}(v) \cdot \mathcal{L}_{psup} \quad (9)$$

where $\mathbb{1}(\cdot)$ is an indicator function that returns 1 if $v \in V_p$, and 0 otherwise. This means that \mathcal{L}_{psup} is considered if the video v is an instance in V_p for which we have the ground truth summary video. The hyperparameters β and γ control the relative importance of the diversity and supervision losses, respectively. We denote this variant of our model as `UnpairedVSNpsup`.

4. Experiments

4.1. Setup

Data and Setting: We conduct evaluation on two standard video summarization datasets: SumMe [10] and TVSum [33]. These datasets have 25 and 50 videos, respectively. Since these datasets are very small, we use another two datasets, namely the YouTube [7] (39 videos) and the OVP dataset [1] (50 videos), to help the learning. Table 1 shows the main characteristics of the datasets. We can observe that these datasets are diverse, especially in terms of ground truth annotations. We follow prior work [8, 39] to convert multiple ground truths with different format to generate a single keyframe-based annotation (a binary key frame indicator vector [39]) for each training video.

From Table 1, we can see that we have in total 164 videos available for experiments. When evaluating on the SumMe dataset, we randomly select 20% of SumMe videos for testing. We use the remaining 80% of SumMe videos and all the videos in other datasets (i.e., TVSum, YouTube and OVP) for training. We create the *unpaired* data from the training subset by first randomly selecting 50% of the raw videos (ignoring their ground truth summaries) and then selecting the ground truth summaries (while ignoring the corresponding raw videos) of the remaining 50% videos. In the end, we obtain a set of raw videos and a set of real summary videos, where there is no correspondence between the raw videos and the summary videos. We follow the same strategy to create the training (unpaired) and testing set when evaluating on the TVSum dataset.

Features: Firstly, we uniformly downsample every video to 2 fps. Then we use *pool5* layer of the pretrained GoogleNet [34] to extract 1024-dimensional feature representation of each frame in the video. Note that our feature extraction follows prior work [24, 31, 39, 47]. This allows us to perform a fair comparison with these works.

Training Details: We train our final model (`UnpairedVSN`) from scratch with a batch size of 1. We use the Adam optimizer [15] with a learning rate of 0.00001 for the key frame selector network (S_K). We use the SGD [3] optimizer with a learning rate of 0.0002 for the summary discriminator network (S_D). We set $\beta = 1$ for SumMe and $\beta = 0.001$ for TVSum in Eq. 6. Additionally, we set β and γ to 0.001 for SumMe and TVSum in Eq. 9.

Evaluation Metrics: We evaluate our method using the keyshot-based metrics as in previous work [24, 39]. Our method predicts summaries in the form of key frames. We convert these key frames to key shots (i.e., an interval-based subset of video frames [10, 11, 39]) following the approach in [39]. The idea is to first temporally segment the videos using KTS algorithm [30]. If a segment contains a key frame, we mark all the frames in that segment as 1, and 0 otherwise. This process may result in many key shots. In or-

Dataset	No. of videos	Content	Ground truth annotation type
SumMe [10]	25	User videos	Interval-based shots and frame-level score
TVSum [33]	50	YouTube videos	Frame-level importance score
YouTube [7] [†]	39	Web videos	Collection of key frames
OVP [1]	50	Various genre videos	Collection of key frames

Table 1. Key characteristics of different datasets used in our experiments. [†]The YouTube dataset has 50 videos, but we exclude (following [8, 39]) the 11 cartoon videos and keep the rest.

der to reduce the number of key shots, we rank the segments according to the ratio between the number of key frames and the length of segment. We then apply knapsack algorithm to generate keyshot-based summaries that are at most 15% of the length of the test video [10, 11, 33, 39]. The SumMe dataset has keyshot-based ground truth annotation, so we directly use it for evaluation. The TVSum dataset provides frame-level importance scores which we also convert to key shots as done by [24, 39] for evaluation.

Given a test video v , let X and Y be the predicted key shot summary and the ground truth summary, respectively. We compute the precision (P), recall (R) and F-score (F) to measure the quality of the summary as follows:

$$P = \frac{\text{overlap in } X \text{ and } Y}{\text{duration of } X}, R = \frac{\text{overlap in } X \text{ and } Y}{\text{duration of } Y} \quad (10)$$

$$F = \frac{2 \times P \times R}{P + R} \quad (11)$$

We follow the evaluation protocol of the datasets (SumMe [10, 11] and TVSum [33]) to compute the F-score between the multiple user created summaries and the predicted summary for each video in the datasets. Following prior work [24], we run our experiments five times for each method and report the average performance over the five runs.

4.2. Baselines

Since our work is the first attempt to learn video summarization using unpaired data, there is no prior work that we can directly compare with. Nevertheless, we define our own baselines as follows:

Unsupervised SUM-FCN: If we remove the summary discriminator network from our model, we can learn video summarization in an unsupervised way. In this case, our learning objective is simply $\mathcal{L}_{reconst} + \mathcal{L}_{div}$. This is equivalent to the unsupervised SUM-FCN in [31]. We call this baseline model SUM-FCN_{unsup}. Note that SUM-FCN_{unsup} is a strong baseline (as shown in [31]) since it already outperforms many existing unsupervised methods ([7, 13, 18, 24, 33, 45]) in the literature.

Model with Adversarial Objective: We define another baseline model where we have the summary discriminator network S_D and the key frame selector network S_K , but the objective to be minimized is $\mathcal{L}_{reconst} + \mathcal{L}_{adv}$ (i.e., we ignore \mathcal{L}_{div}). We refer to this baseline model as UnpairedVSN_{adv}.

4.3. Main Results

	SUM-FCN _{unsup} [31]	UnpairedVSN _{adv}	UnpairedVSN
F-score	44.8	46.5	47.5
Precision	43.9	45.0	46.3
Recall	46.2	49.1	49.4

Table 2. Performance (%) of different methods on the SumMe dataset [10]. We report summarization results in terms of three standard metrics including F-score, Precision and Recall.

	SUM-FCN _{unsup} [31]	UnpairedVSN _{adv}	UnpairedVSN
F-score	53.6	55.3	55.6
Precision	59.1	61.0	61.1
Recall	49.1	50.6	50.9

Table 3. Performance (%) of different methods on TVSum [33].

In Table 2, we provide the results (in terms of F-score, precision and recall) of our final model UnpairedVSN and the baseline models on the SumMe dataset. Our method outperforms the baseline methods on all evaluation metrics. It is also worth noting that when our summary generator and discriminator networks are trained using unpaired data with the adversarial loss (i.e., UnpairedVSN_{adv}), we observe a significant boost in performance (1.7%, 1.1% and 2.9% in terms of F-score, precision and recall, respectively) over the unsupervised baseline SUM-FCN_{unsup}. Adding an additional regularizer \mathcal{L}_{div} (i.e., UnpairedVSN) further improves the summarization performance.

Table 3 shows the performance of different methods on the TVSum dataset. Again, our final method outperforms the baseline methods. Moreover, the trend in performance boost is similar to what we observe on the SumMe dataset.

Results in Table 2 and Table 3 demonstrate that learning from unpaired data is advantageous as it can significantly improve video summarization models over purely unsupervised approaches.

4.4. Comparison with Supervised Methods

We also compare the performance of our method with state-of-the-art supervised methods for video summarization. Recent supervised methods [24, 31, 38, 39, 40, 43, 47] also use additional datasets (i.e., YouTube and OVP) to increase the number of paired training examples while training on the SumMe or the TVSum dataset. For example,

when experimenting on SumMe, they use 20% for testing and use the remaining 80% videos of SumMe along with the videos in TVSum, OVP and YouTube for training. However, the main difference is that we further divide the combined training dataset to create unpaired examples (see Sec. 4.1). In other words, given a pair of videos (a raw video and its summary video), we either keep the raw video or the summary video in our training set. In contrast, both videos are part of the training set in supervised methods. As a result, supervised methods use twice as many videos during training. In addition, supervised methods have access to the correspondence between the raw video and the ground truth summary video. Therefore, it is important to note that the supervised methods utilize far more supervision than our proposed method. We show the comparison in Table 4.

Surprisingly, on the SumMe dataset, our final method outperforms most of the supervised methods (except [31]) by a big margin (nearly 3%). On the TVSum dataset, we achieve slightly lower performance. Our intuition is that if we have more unpaired data for training, we can reduce the performance gap on TVSum. To sum up, this comparison study demonstrates that our unpaired learning formulation has potential to compete with supervised approaches.

Method	SumMe	TVSum
Zhang <i>et al.</i> [38]	41.3	–
Zhang <i>et al.</i> [39] (vsLSTM)	41.6	57.9
Zhang <i>et al.</i> [39] (dppLSTM)	42.9	59.6
Mahasseni <i>et al.</i> [24] (supervised)	43.6	61.2
Zhao <i>et al.</i> [43] [‡]	43.6	61.5
Zhou <i>et al.</i> [47] (supervised)	43.9	59.8
Zhang <i>et al.</i> [40]	44.1	63.9
Rochan <i>et al.</i> [31]	51.1	59.2
UnpairedVSN (Ours)	47.5	55.6

Table 4. Quantitative comparison (in terms of F-score %) between our methods and state-of-the-art supervised methods on SumMe [10] and TVSum [33]. [‡]Results are taken from [40].

4.5. Effect of Partial Supervision

We also examine the performance of our model when direct supervision (i.e., correspondence between videos in V and S) is available for a small number of videos in the training set. Our aim is to study the effect of adding partial supervision to the framework.

In this case, for the first 10% of original/raw videos that are fed to the key frame selector network, we use their ground truth key frame annotations as an additional learning signal (see Eq. 9). Intuitively, we should be able to obtain better performance than learning only with unpaired data, since we have some extra supervision during training.

Table 5 shows the performance of our model trained with this additional partial supervision. We observe a trend

of improvement (across all evaluation metrics) on both the datasets. This shows that our proposed model can be further improved if we have access to some paired data in addition to unpaired data during the training.

	SumMe	TVSum
F-score	48.0 (47.5)	56.1 (55.6)
Precision	46.7 (46.3)	61.7 (61.1)
Recall	49.9 (49.4)	51.4 (50.9)

Table 5. Performance (%) of UnpairedVSN_{psup} on the SumMe [10] and TVSum [33] datasets. In the bracket, we include the performance of our final model UnpairedVSN reported in Table 2 and Table 3 to help with the comparison.

4.6. Transfer Data Setting

In our standard data setting (see Sec. 4.1), it is possible that some of the unpaired examples consist of raw videos or video summaries from the dataset under consideration. In order to avoid this, we conduct additional experiments under a more challenging data setting where the unpaired examples originate totally from different datasets. For instance, if we evaluate on SumMe, we use the videos and user summaries of TVSum, OVP and YouTube to create unpaired training data, and then use the entire SumMe for testing. We follow the similar process while evaluating on TVSum. This kind of data setting is referred as transfer data setting [38, 39], though it has been defined in the context of fully supervised learning. We believe that this data setting is closer to real scenarios, where we may need to summarize videos from domains that are different from those used in training.

Table 6 and Table 7 show the performance of different approaches on SumMe and TVSum, respectively. Although we notice slight degradation in performance compared with the standard data setting, the trend in results is consistent with our findings in Sec. 4.3.

	SUM-FCN _{unsup} [31]	UnpairedVSN _{adv}	UnpairedVSN
F-score	39.5	41.4	41.6
Precision	38.3	40.4	40.5
Recall	41.2	43.6	43.7

Table 6. Performance (%) of different methods on SumMe [10] under transfer data setting.

	SUM-FCN _{unsup} [31]	UnpairedVSN _{adv}	UnpairedVSN
F-score	52.9	55.0	55.7
Precision	58.2	60.6	61.2
Recall	48.5	50.4	51.1

Table 7. Performance (%) of different methods on TVSum [33] under transfer data setting.



Figure 3. Two example results from the SumMe dataset [10]. The two bars at the bottom show the summaries produced by *UnpairedVSN* and humans, respectively. The black bars denote the selected sequences of frames, and the blue bar in background indicate the video length.

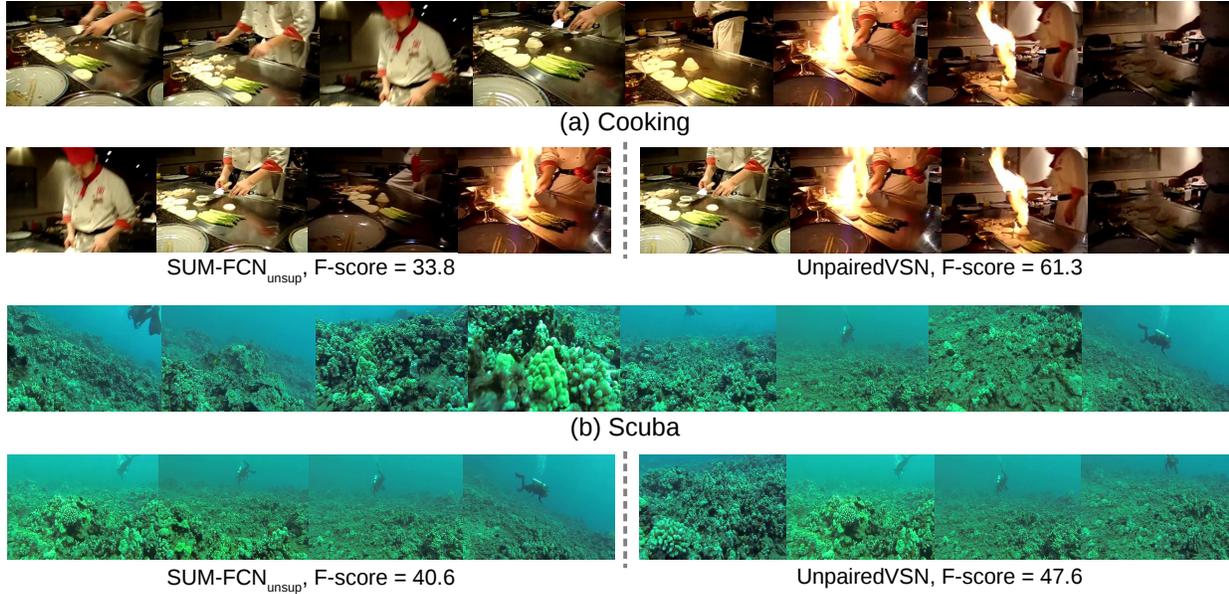


Figure 4. Example videos from SumMe [10] and predicted summaries by *SUM-FCN_{unsup}* [31] and *UnpairedVSN*. Frames in the first row are sampled from the video, whereas frames in the second row are sampled from the summaries generated by different approaches.

4.7. Qualitative Analysis

Figure 3 presents example summaries generated by our method *UnpairedVSN*. We observe that the output summaries from our approach have a higher overlap with the human generated summaries. This implies that our method is able to preserve information essential for generating optimal and meaningful summaries.

We compare the results of different approaches in Fig. 4. The first video in Fig. 4(a) is related to cooking. *SUM-FCN_{unsup}* extracts the shots from the middle of the video and misses the important video shots towards the end. In contrast, we observe that *UnpairedVSN* preserves the temporal story of the video by extracting video shots from different sections while focusing on key scenes. This has resulted in better agreement with the human created summaries. The second video in Fig. 4(b) is about scuba diving. Unlike the first video, there is not a huge performance gap between *SUM-FCN_{unsup}* and *UnpairedVSN*. However, it still noticeable that *SUM-FCN_{unsup}* captures less

diverse scenes compared with *UnpairedVSN*.

5. Conclusion

We have presented a new formulation for video summarization where the goal is to learn video summarization using unpaired training examples. We have introduced a deep learning framework that operates on unpaired data and achieves much better performance than the baselines. Our proposed method obtains results that are even comparable to state-of-the-art supervised methods. If a small number of paired videos are available during training, our proposed framework can be easily extended to take advantage of this extra supervision to further boost the performance. Since unpaired training data are much easier to collect, our work offers a promising direction for future research in video summarization. As future work, we plan to experiment with large-scale unpaired videos collected in the wild.

Acknowledgments: The authors acknowledge financial support from NSERC and UMGF funding. We also thank NVIDIA for donating some of the GPUs used in this work.

References

- [1] Open video project. <https://open-video.org/>.
- [2] Amjad Almahairi, Sai Rajeswar, Alessandro Sordani, Philip Bachman, and Aaron Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. In *International Conference on Machine Learning*, 2018.
- [3] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *International Conference on Computational Statistics*, 2010.
- [4] Sijia Cai, Wangmeng Zuo, Larry S Davis, and Lei Zhang. Weakly-supervised video summarization using variational encoder-decoder and web prior. In *European Conference on Computer Vision*, 2018.
- [5] Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [6] Wen-Sheng Chu, Yale Song, and Alejandro Jaimes. Video co-summarization: Video summarization by visual co-occurrence. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [7] Sandra Eliza Fontes De Avila, Ana Paula Brandão Lopes, Antonio da Luz, and Arnaldo de Albuquerque Araújo. Vsumm: A mechanism designed to produce static video summaries and a novel evaluation method. *Pattern Recognition Letters*, 32(1):56–68, 2011.
- [8] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. Diverse sequential subset selection for supervised video summarization. In *Advances in Neural Information Processing Systems*, 2014.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- [10] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool. Creating summaries from user videos. In *European Conference on Computer Vision*, 2014.
- [11] Michael Gygli, Helmut Grabner, and Luc Van Gool. Video summarization by learning submodular mixtures of objectives. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [12] Hong-Wen Kang and Xue-Quan Chen. Space-time video montage. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2006.
- [13] Aditya Khosla, Raffay Hamid, Chih-Jen Lin, and Neel Sundaresan. Large-scale video summarization using web-image priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [14] Gunhee Kim and Eric P Xing. Reconstructing storyline graphs for image recommendation from web community photos. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [15] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [16] Robert Laganière, Raphael Bacco, Arnaud Hocevar, Patrick Lambert, Grégory Païs, and Bogdan E Ionescu. Video summarization from spatio-temporal features. In *Proceedings of the 2nd ACM TREC Vid Video Summarization Workshop*, 2008.
- [17] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. Discovering important people and objects for egocentric video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [18] Yingbo Li and Bernard Merialdo. Multi-video summarization based on video-mmr. In *Workshop on Image Analysis for Multimedia Interactive Services*, 2010.
- [19] Yandong Li, Liqiang Wang, Tianbao Yang, and Boqing Gong. How local is the local diversity? reinforcing sequential determinantal point processes with dynamic ground sets for supervised video summarization. In *European Conference on Computer Vision*, 2018.
- [20] David Liu, Gang Hua, and Tsuhan Chen. A hierarchical visual model for video object summarization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2178–2190, 2010.
- [21] Tiecheng Liu and John R Kender. Optimization algorithms for the selection of key frame sequences of variable length. In *European Conference on Computer Vision*, 2002.
- [22] Zheng Lu and Kristen Grauman. Story-driven summarization for egocentric video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [23] Yu-Fei Ma, Lie Lu, Hong-Jiang Zhang, and Mingjing Li. A user attention model for video summarization. In *ACM Multimedia*, 2002.
- [24] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. Unsupervised video summarization with adversarial LSTM networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [25] Padmavathi Mundur, Yong Rao, and Yelena Yesha. Keyframe-based video summarization using delaunay clustering. *International Journal on Digital Libraries*, 6(2):219–232, 2006.
- [26] Jeho Nam and Ahmed H Tewfik. Event-driven video abstraction and visualization. *Multimedia Tools and Applications*, 16(1-2):55–77, 2002.
- [27] Chong-Wah Ngo, Yu-Fei Ma, and Hong-Jiang Zhang. Automatic video summarization by graph modeling. In *IEEE International Conference on Computer Vision*, 2003.
- [28] Rameswar Panda, Abir Das, Ziyang Wu, Jan Ernst, and Amit K Roy-Chowdhury. Weakly supervised summarization of web videos. In *IEEE International Conference on Computer Vision*, 2017.
- [29] Rameswar Panda and Amit K Roy-Chowdhury. Collaborative summarization of topic-related videos. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [30] Danila Potapov, Matthijs Douze, Zaid Harchaoui, and Cordelia Schmid. Category-specific video summarization. In *European Conference on Computer Vision*, 2014.
- [31] Mrigank Rochan, Linwei Ye, and Yang Wang. Video summarization using fully convolutional sequence networks. In *European Conference on Computer Vision*, 2018.

- [32] Aidean Sharghi, Ali Borji, Chengtao Li, Tianbao Yang, and Boqing Gong. Improving sequential determinantal point processes for supervised video summarization. In *European Conference on Computer Vision*, 2018.
- [33] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. Tvsum: Summarizing web videos using titles. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [35] Ziyou Xiong, Regunathan Radhakrishnan, Ajay Divakaran, Zou Yong-Rui, and Thomas S Huang. *A unified framework for video summarization, browsing & retrieval: with applications to consumer and surveillance video*. Elsevier, 2006.
- [36] Zili Yi, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *IEEE International Conference on Computer Vision*, 2017.
- [37] Hong Jiang Zhang, Jianhua Wu, Di Zhong, and Stephen W Smoliar. An integrated system for content-based video retrieval and browsing. *Pattern recognition*, 30(4):643–658, 1997.
- [38] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Summary transfer: Exemplar-based subset selection for video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [39] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Video summarization with long short-term memory. In *European Conference on Computer Vision*, 2016.
- [40] Ke Zhang, Kristen Grauman, and Fei Sha. Retrospective encoders for video summarization. In *European Conference on Computer Vision*, 2018.
- [41] Yujia Zhang, Michael Kampffmeyer, Xiaodan Liang, Min Tan, and Eric P Xing. Query-conditioned three-player adversarial network for video summarization. In *British Machine Vision Conference*, 2018.
- [42] Yujia Zhang, Michael Kampffmeyer, Xiaodan Liang, Dingwen Zhang, Min Tan, and Eric P Xing. Dtr-gan: Dilated temporal relational adversarial network for video summarization. *arXiv preprint arXiv:1804.11228*, 2018.
- [43] Bin Zhao, Xuelong Li, and Xiaoqiang Lu. Hierarchical recurrent neural network for video summarization. In *ACM Multimedia*, 2017.
- [44] Bin Zhao, Xuelong Li, and Xiaoqiang Lu. Hsa-rnn: Hierarchical structure-adaptive rnn for video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [45] Bin Zhao and Eric P Xing. Quasi real-time summarization for consumer videos. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [46] Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network. In *International Conference on Learning Representations*, 2017.
- [47] Kaiyang Zhou, Yu Qiao, and Tao Xiang. Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. In *AAAI Conference on Artificial Intelligence*, 2018.
- [48] Kaiyang Zhou, Tao Xiang, and Andrea Cavallaro. Video summarisation by classification with deep reinforcement learning. In *British Machine Vision Conference*, 2018.
- [49] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision*, 2017.