

Rethinking the Evaluation of Video Summaries

Mayu Otani
CyberAgent, Inc.

Yuta Nakashima
Osaka University

Esa Rahtu
Tampere University

Janne Heikkilä
University of Oulu

Abstract

Video summarization is a technique to create a short skim of the original video while preserving the main stories/content. There exists a substantial interest in automating this process due to the rapid growth of the available material. The recent progress has been facilitated by public benchmark datasets, which enable easy and fair comparison of methods. Currently the established evaluation protocol is to compare the generated summary with respect to a set of reference summaries provided by the dataset. In this paper, we will provide in-depth assessment of this pipeline using two popular benchmark datasets. Surprisingly, we observe that randomly generated summaries achieve comparable or better performance to the state-of-the-art. In some cases, the random summaries outperform even the human generated summaries in leave-one-out experiments. Moreover, it turns out that the video segmentation, which is often considered as a fixed pre-processing method, has the most significant impact on the performance measure. Based on our observations, we propose alternative approaches for assessing the importance scores as well as an intuitive visualization of correlation between the estimated scoring and human annotations.

1. Introduction

The tremendous growth of the available video material has escalated the demand for techniques that enable users to quickly browse and watch videos. One remedy is provided by the automatic video summarization, where the aim is to produce a short video skim that preserve the most important content of the original video. For instance, the original footage from a sport event could be compressed into a few minute summary illustrating the most important events such as goals, penalty kicks, etc.

Numerous automatic summarization methods have been proposed in the literature. The most recent methods follow a paradigm that consists of a video segmentation, importance score prediction, and video segment selection as illustrated in Figure 1. The most challenging part of this pipeline is the importance score prediction, where the task is to highlight

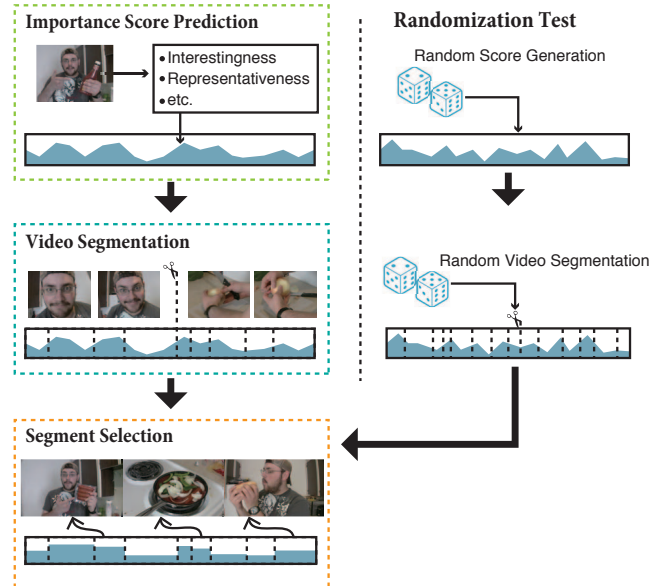


Figure 1. An illustration of the commonly used video summarization pipeline and our randomization test. We utilize random summaries to validate the current evaluation frameworks.

the parts that are most important for the video content. Various factors affect importance of video parts, and different video summaries are possible for a single video given a different criterion of importance. In fact, previous works have proposed a variety of importance criteria, such as visual interestingness [2, 3], compactness (*i.e.*, smaller redundancy) [28], and diversity [25, 29].

Despite the extensive efforts toward automatic video summarization, the evaluation of the generated summaries is still an unsolved problem. A straightforward but yet convincing approach would be to utilise a subjective evaluation; however, collecting human responses is expensive, and reproduction of the result is almost impossible due to the subjectivity. Another approach is to compare generated video summaries to a set of fixed reference summaries prepared by human annotators. To this end, the human annotators are asked to create video summaries, which are then treated as ground truth. The advantage of this approach is the reusability of the reference summaries, *i.e.*, different video summarization methods can be evaluated without additional anno-

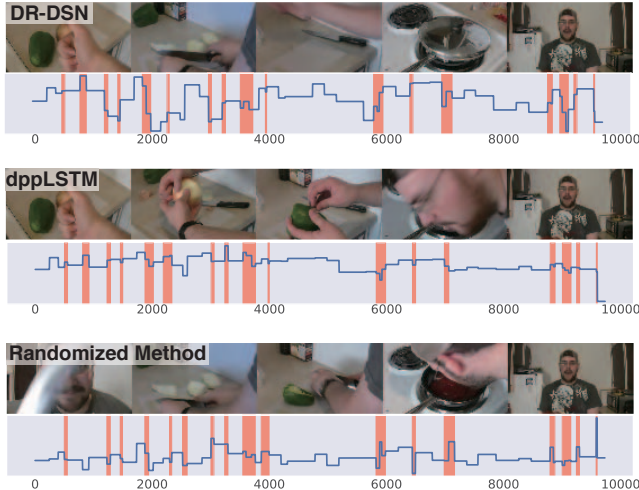


Figure 2. Comparison of summaries created by two recent methods and our randomized method (Section 4). The blue line shows the segment level importance scores with respect to time (frames). The orange areas indicate the frames selected for the final summary. All of the three methods use the same segment boundaries by KTS [14]. Interestingly, all methods (including the random one) produce very similar outputs despite clear differences in the importance scores.

tations and the experiments can be reproduced.

The most popular datasets used for reference based evaluations are SumMe [2] and TVSum [18]. These datasets provide a set of videos as well as multiple human generated reference summaries (or importance scores) for each original video. The basic evaluation approach, used with both datasets, is to measure the agreement between the generated summary and the reference summaries using F1 score. Since their introduction, SumMe and TVSum have been widely adopted in the recent video summarization literature [3, 12, 21, 25, 26, 27, 29]. Nevertheless, the validity of reference summary-based evaluation has not been previously investigated.

This paper delves deeper into the current reference based evaluation framework using SumMe [2] and TVSum [18] datasets. We will first review the framework and then apply a randomization test to assess the quality of the results. The proposed randomization test generates video summaries based on random importance scores and random video segmentation. Such summaries provide a baseline score that is achievable by chance.

Figure 2 illustrates one of the main findings of our work. It turned out that the random method produces summaries that are almost identical to the state-of-the-art despite the fact that it is not using the video content at all for importance score prediction. Deeper analysis shows that while there are differences in the importance scores they are ignored when assembling the final summary. The randomization test revealed critical issues in the current video summa-

rization evaluation protocol, which motivated us to propose a new framework for assessing the importance rankings.

The main contributions of this paper are as follows:

- We assess the validity of the current reference summary-based evaluation framework and reveal that a random method is able to reach similar performance scores as the current state-of-the-art.
- We demonstrate that the widely used F1 score is mostly determined by the distribution of video segment lengths. Our analysis provides a simple explanation for this phenomenon.
- We demonstrate evaluating the importance rankings using correlation between the predicted ordering and the ordering by human annotators. Moreover, we propose several visualisations that give insight to the predicted scoring versus random scores.

2. Related Work

2.1. Video Summarization

A diverse set of video summarization approaches have been presented in the literature. One group of works aim at detecting important shots by measuring the visual interestingness [2], such as dynamics of visual features [8], and visual saliency [11]. Gygli *et al.* [3] combined multiple properties including saliency, aesthetics, and presence of people in the frames.

Another group of methods aims at compactness by discarding redundant shots [28]. Maximization of representativeness and diversity in the output video are also widely used criteria in the recent works [1, 14, 25]. These methods are based on the assumption that a good summary should have diverse content while the sampled shots explain the events in the original video.

Recently, LSTM-based deep neural network models have been proposed to directly predict the importance scores given by the human annotators [26]. The model is also extended with determinantal point process [7] to ensure diverse segment selection. Finally, Zhou *et al.* [21] applied reinforcement learning to obtain a policy for the frame selection in order to maximize the diversity and representativeness of the generated summary.

Although these works use various importance criteria, many of them employ a similar processing pipeline. Firstly, the importance scores are produced for each frame in the original video. Secondly, the obtained video is divided into short segments. Finally, the output summary is generated by selecting a subset of video segments by maximising the importance scores with the knapsack constraint.

Table 1. The F1 measures for SumMe and TVSum benchmarks as reported in recent works. Average (Avr) denotes the average of F1 scores over all reference summaries and maximum (Max) denotes the highest F1 score within the reference summaries [3]. In addition, we show the F1 values for our randomized test and human annotations (leave-one-out test). It can be noted that random summaries achieve comparable results to the state-of-the-art and even to human annotations.

	Video segmentation	SumMe		TVSum	
		Avr.	Max.	Avr.	Max.
LSVS [5]		—	—	0.36	—
QRTS [28]	Uniform segmentation	—	—	0.46	—
CSUV [2]	Superframe segmentation	0.23	—	—	—
TVSum [18]	Change-point detection	—	—	0.50	—
VS-LMM [3]	Uniform segmentation	—	0.40	—	—
dppLSTM [26]	KTS	—	0.43	0.60	—
VS-DSF [12]	Uniform segmentation	0.18	—	—	—
Summary Transfer [25]	KTS	—	0.41	—	—
DR-DSN [29]	KTS	—	0.41	0.58	—
re-seq2seq [27]	LSTM-based segmentation	—	0.45	0.64	—
SASUM [21]	KTS	—	0.45	0.58	—
Randomized test	KTS	0.19	0.41	0.57	0.78
Randomized test	Two-peak	0.14	0.27	0.58	0.71
Human	KTS	0.31	0.54	0.54	0.78

2.2. Video Summary Evaluation

The evaluation of a video summary is a challenging task. This is mainly due to subjective nature of the quality criterion that varies from viewer to viewer and from one time instant to another. The limited number of evaluation videos and annotations further magnify this ambiguity problem.

Most early works [10, 11, 19] as well as some recent works [22] employ user studies, in which viewers subjectively score the quality of output video summaries prepared solely for the respective works [10, 15, 23]. The critical shortcoming in such approach is the related cost and reproducibility. That is, one cannot obtain the same evaluation results, even if the same set of viewers would re-evaluate the same videos.

Many recent works instead evaluate their summaries by comparing them to reference summaries. Khosla *et al.* [5] proposed to use the pixel-level distance between keyframes in reference and generated summaries. Lee *et al.* [9] use number of frames that contain objects of interest as a similarity measure. Gong *et al.* [1] compute precision and recall scores over keyframes selected by human annotators. Yeung *et al.* [24] propose a different approach and evaluate the semantic similarity of the summaries based on textual descriptions. For this, they generated a dataset with long egocentric videos for which the segments are annotated with textual descriptions. This framework is mainly used to evaluate video summaries based on user queries [13, 16]. More recently, computing overlap between reference and generated summaries has become the standard framework for video summary evaluation [2, 3, 14, 17, 18, 28].

This paper investigates the evaluation framework where

generated summaries are compared to a set of human annotated references. Currently, there are two public datasets that facilitate this type of evaluation. SumMe [2] and TVSum [18] datasets provide manually created reference summaries and are currently the most popular evaluation benchmarks. The SumMe dataset contains personal videos and the corresponding reference summaries collected from 15–18 annotators. The TVSum dataset provides shot-level importance scores for YouTube videos. Most of the literature uses the F1 measure between generated summaries and reference summaries as a performance indicator. Table 1 shows reported scores for both datasets. The SumMe dataset, which has around 15 different reference summaries, has two possible ways for aggregating the F1 scores: One is to compute an average of F1 measures over all reference summaries, and the other is to use the maximum score.

3. Current evaluation framework

3.1. SumMe

SumMe is a video summarization dataset that contains 25 personal videos obtained from the YouTube. The videos are unedited or minimally edited. The dataset provides 15–18 reference summaries for each video. Human annotators individually made the reference summaries so that the length of each summary is less than 15% of the original video length. For evaluation, generated summaries should be subject to the same constraint on the summary length.

3.2. TVSum

TVSum contains 50 YouTube videos, each of which has a title and a category label as metadata. Instead of provid-

ing reference summaries, the TVSum dataset provides human annotated importance scores for every two second of each video. For evaluation, the reference summaries, with a predefined length, are generated from these importance scores using the following procedure: Firstly, videos are divided into short video segments, which are the same as in the generated summary. Then, the importance scores within a video segment are averaged to obtain a segment-level importance score. Finally, a reference summary is generated by finding a subset of segments that maximizes the total importance score in the summary. The advantage of this approach is the ability to generate summaries with desired length.

3.3. Evaluation measure

The most common evaluation approach is to compute F1 measure between the predicted and the reference summaries. Let $y_i \in \{0, 1\}$ denote a label indicating which frames from the original video is selected to the summary (*i.e.* $y_i = 1$ if the i -th frame is selected and otherwise 0). Given similar label y_i^* for the references summary, the F1 score is defined as

$$F1 = \frac{2PRE \cdot REC}{PRE + REC}, \quad (1)$$

where

$$PRE = \frac{\sum_{i=1}^N y_i \cdot y_i^*}{\sum_{i=1}^N y_i} \quad \text{and} \quad REC = \frac{\sum_{i=1}^N y_i \cdot y_i^*}{\sum_{i=1}^N y_i^*}, \quad (2)$$

are the frame level precision and recall scores. N denotes the total number of frames in the original video.

In the experiments, the F1 score is computed for each reference summary separately and the scores are summarised either by averaging or selecting the maximum for each video. The former approach implies that the generated summary should include segments with largest number of agreement, while the latter argue that all human annotators provided reasonable importance scores and thus the generated summary should have high score if it matches at least one of the reference summaries.

4. Randomization test

Commonly video summarization pipeline consists of three components; importance score estimation, video segmentation, and shot selection (Figure 1). We devise a randomization test to evaluate the contribution of each part to the final evaluation score. In these experiments we generate video summaries that are independent of video content by utilising random importance scores and random video segment boundaries. Specifically, the importance score for each frame is drawn independently from an uniform distribution $[0, 1]$. When needed, the segment-level scores are

produced by average pooling the corresponding frame-level random scores. For video segmentation, we utilise the options defined below.

Uniform segmentation divides the video into segments of constant duration. We used 60 frames in our experiments, which roughly corresponds to 2 seconds (the frame rates in SumMe and TVSum datasets are 30 fps and 25 fps, respectively).

One-peak segmentation samples the number of frames in each segment from an unimodal distribution. We assume that the number of frames between adjacent shot boundaries follow the Poisson distribution with event rate $\lambda = 60$.

Two-peak segmentation is similar to one-peak version, but utilises bimodal distribution, *i.e.*, a mixture of two Poisson distributions, whose event rates are $\lambda = 30$ and $\lambda = 90$, respectively. For sampling, we randomly choose one of the two Poisson distributions with the equal probability and then sample the number of frames. Consequently, a video is segmented into both longer and shorter segments, yet the expected number of frames in a segment is 60 frames.

In addition to the completely random methods, we assess one commonly used segmentation approach and its variation in conjunction with the random scores.

Kernel temporal segmentation (KTS) [14] is based on the visual content of a video and is the most widely used method in the recent video summarization literature (Table 1). KTS produces segment boundaries by detecting changes in visual features. A video segment tends to be long if visual features do not change considerably.

Randomized KTS first segments the video with KTS and then shuffles the segment ordering; therefore, the distribution of segment lengths is exactly the same as KTS's, but the segment boundaries are not synchronized with the visual features.

F1 scores obtained by these randomized (and partially randomized) summaries serve as a baseline that can be achieved completely by chance. Reasonable evaluation framework should produce higher scores for methods that are producing sensible importance scores. Furthermore, one would expect that human generated ground truth summaries should produce top scores in leave one out experiments.

4.1. Analysis on the SumMe dataset

Figure 3 displays the F1 scores (average and maximum) obtained with different versions of our randomized method (see previous section). We performed 100 trials for every random setting and the black bar is the 95% confidence interval. In addition, the same figure contains the corresponding F1 scores for each random segmentation method, but using frame level importance scores from one recently published methods DR-DSN [29]. The reference performance is obtained using human created reference summaries in leave-one-out scheme. In this case, the final result is cal-




Figure 3. F1 scores for different segmentation and importance score combinations for SumMe. Light blue bars refer to random summaries and dark blue bars indicate scores of manually created reference summaries (leave-one-out test). Purple bars show the scores for DR-DSN importance scoring with different segmentation methods. Left: the average of mean F1 scores over reference summaries. Right: the average of the maximum scores.




Figure 4. Recently reported F1 scores for methods using KTS segmentation in SumMe. The average score for random summaries with KTS segmentation is represented by a light blue dashed line.

culated by averaging the F1 scores (avg or max) obtained for each reference summary.

Interestingly, we observe that the performance is clearly dictated by the segmentation method and there is small (if any) impact on the importance scoring. Moreover, the difference between human performance and the best performing automatic method is similar in magnitude to the differences between the segmentation approaches. Figure 4 illustrates the recent state-of-the-art results for SumMe dataset. Surprisingly, KTS segmentation with random importance scores obtains comparable performance to the best published methods. Section 4.3 provides possible explanations for this phenomenon.

4.1.1 Human Evaluation on SumMe

We conducted human evaluation to compare summaries on the SumMe dataset. Subjects compare two video summaries and determine which video better summarizes the original video. In the first experiment, we asked subjects to rate video summaries generated using random importance scores and DR-DSN scores. Both methods use KTS segmentation. Overall, random scores got a slightly higher score than DR-DSN, however, 46% of answers were that the summaries are equally good (bad). This result agree with the observation in the Section 4.1 that the importance scoring hardly affects the evaluation score on the SumMe dataset. We also compare KTS and uniform segmentation with random importance scoring. As a result, subjects prefer uniform segmentation for videos recording long activity, *e.g.*, sightseeing of the statue of liberty and scuba diving. On the other hand, KTS works better for videos with notable events or activities. For such videos, the important parts have little ambiguity, therefore the F1 scores based on the agreement between generated summaries and reference summaries can get higher. For the detailed results of the human evaluation, see the supplementary material.

4.2. Analysis on TVSum dataset

Instead of reference summaries, TVSum dataset contains human annotated importance scores for every 2 second segment in the original video. The main advantage of this approach is the ability to generate reference summaries of arbitrary length. It is also possible to use different segmentation methods. For these reasons, TVSum provides an excellent tool for studying the role of importance scoring and segmentation in the current evaluation framework.

Figure 5 displays the F1 scores for different segmentation methods using both random and the human annotated importance scores. In the latter case, the results are computed using leave-one-out procedure. Surprisingly, for the most of the segmentation methods, the random importance scores have similar performance as human annotations. In addition, the completely random two-peak segmentation performs equally well as content based KTS segmentation. Furthermore, the results in Table 1 illustrate that our random results are on-par (or at least comparable) with the best reported results in the literature. The uniform and one-peak segmentation do not reach the same results, but in these cases the better importance scoring seems to help. In general, the obtained results highlight the challenges in utilizing the current F1 based evaluation frameworks.

4.3. Discussion

As observed in the previous sections, the random summaries resulted in surprisingly high performance scores. The results were on-par with the state-of-the-art and sometimes surpassed even the human level scores. In particular, the segmentation methods that produce large variation in the segment length (*i.e.* two-peak, KTS, and randomized KTS) produced high F1 scores. The results may be understood by examining how the segment length affects on selection procedure in the knapsack formulation that is most commonly adopted in video summarization methods.




Figure 5. F1 scores for different segmentation methods combined to either random or human annotated importance scores (leave-one-out) for TVSum dataset. Light blue bars refer to random scores and dark blue bars indicates human annotations. Interestingly, the random and human annotations obtain similar F1 scores in most cases.

A dynamic programming solver, commonly used for the knapsack problem, selects a segment only if the corresponding effect on the overall score is larger than that of any combination of remaining segments whose total length is shorter. In other words, a segment A is selected only if there are no segments B and C whose combined length is less than A and the effect to total score is more or equal to A. This is rarely true for longer segments in the current summarization tasks, and therefore the summary is usually only composed of short segments. This phenomenon significantly limits the reasonable choices available for segment subset selection. For example, two-peak segmentation draws a segment length from two distributions whose modes are 30 frames and 90 frames; therefore, we can roughly say that longer segments occupies two-third of the total length. If these longer segments are all discarded, the generated summary only consists of the rest one-third of the original video. For generating a summary whose length is 15% of the original video duration, most of the segments are expected to be shared for generated and reference summaries regardless of associated importance scores. This is illustrated in Figure 6. Due to the same reason, the importance scores have more impact if all the segments have equal length (see uniform and one-peak results in Figure 5).

Using the sum of frame-level scores may alleviate the challenge; however, most works instead employ averaging because this drastically increases F1 scores on TVSum. With summation, human summary clearly outperforms random ones, but we can still see the effect of segmentation.

The results on SumMe dataset in Section 4.1 illustrate another challenge. For this dataset, KTS-based references obtain really high performance scores. The use of KTS implicitly incorporate small-redundancy strategy, which aims to create a visually non-redundant video summary. That is, KTS groups visually similar frames into a single segment. Therefore, long segments are likely to be redundant and less lively and thus they are less interesting. Human annotators would not like to include such segments in their

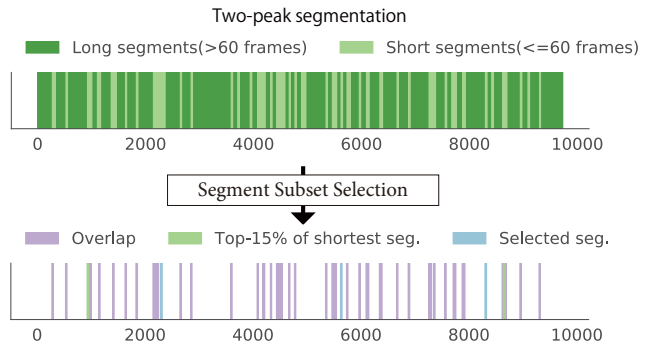


Figure 6. Long segments are implicitly discarded from the summary and only short segments are selected. Top: Green and light green areas visualize segment boundaries generated by the two-peak segmentation method. Bottom plot shows segments selected by dynamic programming algorithm (blue), and top 15% of the shortest video segments (light green), and segments overlapping between them (Purple). Notice that the most of the selected parts are within the group of the shortest segments.

summaries. Meanwhile, the dynamic programming-based segment subset selection tends to avoid long segments as discussed above. Thus generated summaries tend to match the human preference.

5. Importance score evaluation framework

The aforementioned challenges render the current benchmarks inapplicable for assessing the quality of the importance scores. At the same time, most of the recent video summarization literature present methods particularly for importance score prediction. To overcome this problem, we present a new alternative approach for the evaluation.

5.1. Evaluation using rank order statistics

In statistics, rank correlation coefficients are well established tools for comparing the ordinal association (i.e. relationship between rankings). We take advantage of these tools in measuring the similarities between the implicit

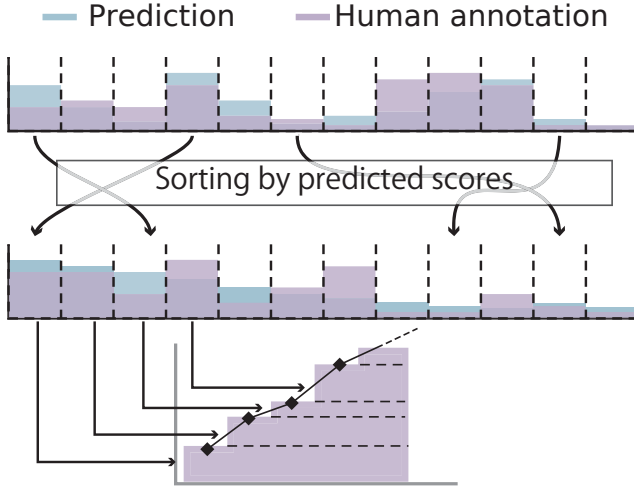


Figure 7. Overview of the score curve formation.

rankings provided by generated and human annotated frame level importance scores as in [20].

More precisely, we use Kendall’s τ [4] and Spearman’s ρ [6] correlation coefficients. To obtain the results, we first rank the video frames according to the generated importance scores and the human annotated reference scores (one ranking for each annotator). In the second stage, we compare the generated ranking with respect to each reference ranking. The final correlation score is then obtained by averaging over the individual results.

We demonstrate the rank order correlation measures, by evaluating two recent video summarization methods (dppLSTM [26] and DR-DSN [29]). For both methods, we utilise the implementations provided by the original authors. For sanity check, we also compute the results using random scoring, which by definition should produce zero average score. These results are obtained by generating 100 uniformly-distributed random value sequences in $[0, 1]$ for each original video and averaging over the obtained correlation coefficients. The human performance is produced using leave-one-out approach. Table 2 summarizes the obtained results for TVSum dataset.

Overall, the metric shows a clear difference between tested methods and the random scoring. In addition, the correlation coefficient for human-annotations is significantly higher than for any other method, which confirms that human importance scores correlate to each other. From the tested methods, dppLSTM results in higher performance compared to DR-DSN. This makes sense, since dppLSTM is particularly trained to predict human annotated importance scores, while DR-DSN aims at maximizing the diversity of the content in the generated summaries. However, both methods clearly outperform the random scoring.

We further investigate the relation between the correlation measures for importance scores and the quality of out-

Table 2. Kendall’s τ and Spearman’s ρ correlation coefficients computed between different importance scores and manually annotated scores on TVSum dataset.

Method	Kendall’s τ	Spearman’s ρ
dppLSTM [26]	0.042	0.055
DR-DSN [29]	0.020	0.026
Random	0.000	0.000
Human	0.177	0.204

put video summaries. We compare video summaries generated using importance scores which positively correlate with human annotations and those using importance scores with negative correlation. The result of human evaluation demonstrated that video summaries generated using importance scores with positive correlation outperformed others. The details of the result are in the supplementary material.

5.2. Visualizing importance score correlations

One of the main challenges in the evaluation of video summaries is the inconsistency between the human-annotations. In fact, although the human annotations result in the highest correlation coefficient in Table 2, the absolute value of the correlation is still relatively low. This stems from subjectivity and ambiguity in the importance score annotation. As we can imagine, what is important in a video can be highly subjective, and the annotators may or may not agree. Furthermore, even if the annotators agree that a certain video content is important, there can be multiple parts in a video that contain the same content in different viewpoints and expressions. Selection from these parts may still be ambiguous problem.

To highlight the variation in the annotations, we propose to visualize the predicted importance score ranking with respect to the reference annotations. To do this, we first compute the frame level average scores over the human annotators. In the second stage, we sort the frames with respect to the predicted importance scores in descending order (Figure 7, middle). In the final step, we accumulate the averaged reference scores based on the ranking obtained in the second stage. More precisely,

$$a_i = \sum_{t=1}^i \frac{s_t}{\sum_{j=1}^n s_j}, \quad (3)$$

where s_i denotes the average human-annotated score for the i -th frame in the sorted video. The normalization factor in the denominator ensures that the maximum value equals to 1. As shown in Figure 7 (bottom), a_i forms a monotonically increasing curve over the sorted frames. If the predicted scores have high correlation to human scores, the curve should increase rapidly. Similar curves can be produced for the human scores using leave-one-out approach.

Figure 8 shows correlation curves produced for two videos from TVSum dataset. The red lines show the a_t

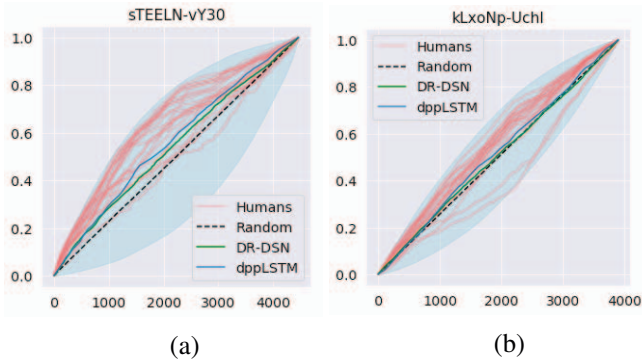


Figure 8. Example correlation curves produced for two videos from TVSum dataset (sTEELN-vY30 and kLxoNp-Uchl are video ids). The red lines represent correlation curves for each human annotator and the black dashed line is the expectation for a random importance scores. The blue and green curves show the corresponding results to dppLSTM and DR-DSN methods, respectively. See supplementary material for more results.

curve for each human annotator and the black dashed line is the expectation for a random importance scores. The blue and green curves show the corresponding results to dppLSTM and DR-DSN methods, respectively. The light-blue colour illustrates the area, where correlation curves may lie. That is, when the predicted importance scores are perfectly concordant with averaged human-annotated scores, *i.e.*, the score based rankings are the same, the curve lies on the upper bound of the light-blue area. On the other hand, a curve coincides with the lower bound of the area when the ranking of the scores is in a reverse order of the reference.

The most of the human annotators obtain a curve that is well above the random baseline in Figure 8. Moreover, Figure 8 (a) shows that both dppLSTM and DR-DSN are able to predict importance scores that are positively correlated with human annotations. On the other hand, Figure 8 (b) shows two red lines that are well below the black dashed line. This implies that these annotators labelled almost opposite responses to the overall consensus. Detailed observation in Figure 9 reveals that this is indeed the case. The outliers highlighted segments around 1500 and 3000 frames, on the other hand, other annotators showed almost opposite opinion for the segments. The proposed visualization provides intuitive tool for illustrating such tendencies.

6. Conclusion

Public benchmark datasets play an important role as they facilitating easy and fair comparison of methods. The quality of the benchmark evaluations have high impact as the research work is often steered to maximise the benchmark results. In this paper, we have assessed the validity of two widely used video summarization benchmarks. Our analysis reveals that the current F1 score based evaluation framework has severe problems.

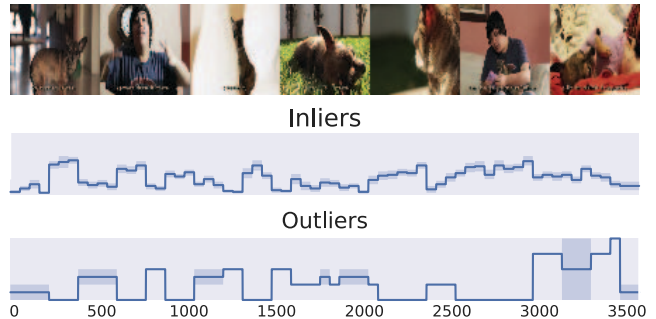


Figure 9. Comparison of human-annotated scores. The bottom row shows the frame level importance scores for the selected two human annotators (outliers). The middle row displays the similar score obtained by averaging over the remaining human annotators (inliers). The top row illustrates keyframes from the corresponding video. One can notice that inliers and outlier have highlighted almost completely opposite parts of the video.

In most cases it turned out that randomly generated summaries were able to reach similar or even better performance scores than the state-of-the-art methods. Sometimes the performance of completely random method surpassed that of human annotators. Closer analysis revealed that score formation was mainly dictated by the video segmentation and particularly the distribution of the segment lengths. This was mainly due to the widely used subset selection procedure. In most cases, the contribution of the importance scores were completely ignored by the benchmark tests.

Based on our observations, we proposed to evaluate the methods using the correlation between predicted and human-annotated importance scores instead of the final summary given by the segment subset selection process. The introduced evaluation offers additional insights about the behaviour of the summarization methods. We also proposed to visualize the correlations by accumulative score curve, which intuitively illustrates the quality of the importance scores with respect to various human annotations.

The proposed new evaluation framework covers only methods that estimate the frame level importance scores. It is not suitable for other approaches such as *e.g.*, clustering-based methods that pick out video segments close to cluster centres. In addition, we primarily addressed the evaluation based on correlation with human annotations. Other factors like comprehensibility of a story in a video, visual aesthetics and relevance to a user query would also be valuable for various applications. We believe that it would be important to address these aspects in future works. Moreover, we believe that new substantially larger datasets are needed for pushing video summarization research forward.

Acknowledgement This work was partly supported by JSPS KAKENHI Grant Nos. 16K16086 and 18H03264.

References

- [1] B. Gong, W.-L. Chao, K. Grauman, and F. Sha. Diverse sequential subset selection for supervised video summarization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2069–2077, 2014.
- [2] M. Gygli, H. Grabner, H. Riemenschneider, and L. van Gool. Creating summaries from user videos. In *European Conference on Computer Vision (ECCV)*, pages 505–520, 2014.
- [3] M. Gygli, H. Grabner, and L. Van Gool. Video summarization by learning submodular mixtures of objectives. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3090–3098, 2015.
- [4] M. G. Kendall. The treatment of ties in ranking problems. *Biometrika*, 33(3):239–251, 1945.
- [5] A. Khosla, R. Hamid, C.-J. Lin, and N. Sundaresan. Large-scale video summarization using web-image priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2698–2705, 2013.
- [6] S. Kokoska and D. Zwillinger. *CRC standard probability and statistics tables and formulae*. Crc Press, 1999.
- [7] A. Kulesza and B. Taskar. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5(2–3), 2012.
- [8] R. Laganière, R. Bacco, A. Hocevar, P. Lambert, G. Païs, and B. E. Ionescu. Video summarization from spatio-temporal features. In *ACM TRECVID Video Summarization Workshop*, pages 144–148, 2008.
- [9] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1346–1353, 2012.
- [10] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2714–2721, 2013.
- [11] Y. Ma, L. Lu, H. Zhang, and M. Li. A user attention model for video summarization. In *ACM International Conference on Multimedia (MM)*, pages 533–542, 2002.
- [12] M. Otani, Y. Nakashima, E. Rahtu, J. Heikkilä, and N. Yokoya. Video summarization using deep semantic features. In *Asian Conference on Computer Vision (ACCV)*, volume 10115, pages 361–377, 2016.
- [13] B. Plummer, M. Brown, and S. Lazebnik. Enhancing video summarization via vision-language embedding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5781–5789, 2017.
- [14] D. Potapov, M. Douze, Z. Harchaoui, and C. Schmid. Category-specific video summarization. In *European Conference on Computer Vision (ECCV)*, pages 540–555, 2014.
- [15] J. Sang and C. Xu. Character-based movie summarization. In *ACM International Conference on Multimedia (MM)*, pages 855–858, 2010.
- [16] A. Sharghi, B. Gong, and M. Shah. Query-focused extractive video summarization. In *European Conference on Computer Vision (ECCV)*, pages 3–19, 2016.
- [17] Y. Song. Real-time video highlights for yahoo esports. In *Neural Information Processing Systems (NIPS) Workshops*, 5 pages, 2016.
- [18] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes. TVSum: Summarizing web videos using titles. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5179–5187, 2015.
- [19] C. M. Taskiran, Z. Pizlo, A. Amir, D. Ponceleon, and E. E. J. Delp. Automated video program summarization using speech transcripts. *IEEE Transactions on Multimedia*, 8(4):775–790, 2006.
- [20] A. B. Vasudevan, M. Gygli, A. Volokitin, and L. Van Gool. Query-adaptive video summarization via quality-aware relevance estimation. In *ACM International Conference on Multimedia (MM)*, pages 582–590, 2017.
- [21] H. Wei, B. Ni, Y. Yan, H. Yu, X. Yang, and C. Yao. Video Summarization via Semantic Attended Networks. In *AAAI Conference on Artificial Intelligence*, pages 216–223, 2018.
- [22] T. Yao, T. Mei, and Y. Rui. Highlight detection with pairwise deep ranking for first-person video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] T. Yao, T. Mei, and Y. Rui. Highlight detection with pairwise deep ranking for first-person video summarization. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] S. Yeung, A. Fathi, and L. Fei-fei. VideoSET : Video summary evaluation through text. *arXiv preprint arXiv:1406.5824v1*, 2014.
- [25] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Summary transfer: Exemplar-based subset selection for video summarization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1059–1067, 2016.
- [26] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Video summarization with long short-term memory. In *European Conference on Computer Vision (ECCV)*, pages 766–782, may 2016.
- [27] K. Zhang, K. Grauman, and F. Sha. Retrospective Encoders for Video Summarization. In *European Conference on Computer Vision (ECCV)*, pages 383–399, 2018.
- [28] B. Zhao and E. P. Xing. Quasi real-time summarization for consumer videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2513–2520, 2014.
- [29] K. Zhou, Y. Qiao, and T. Xiang. Deep reinforcement learning for unsupervised video summarization with diversity-representativeness reward. 2018.