

# Show and Recall: Learning What Makes Videos Memorable

Sumit Shekhar  
Adobe Research

sushekha@adobe.com

Dhruv Singal  
Adobe Research

dhsingal@adobe.com

Harvineet Singh  
Adobe Research

harvines@adobe.com

Akhil Shetty  
University of California, Berkeley

shetty.akhil@berkeley.edu

Manav Kedia  
University of Illinois at Urbana-Champaign

mkedial2@illinois.edu

## Abstract

*With the explosion of video content on the Internet, there is a need for research on methods for video analysis which take human cognition into account. One such cognitive measure is memorability, or the ability to recall visual content after watching it. Prior research has looked into image memorability and shown that it is intrinsic to visual content, but the problem of modeling video memorability has not been addressed sufficiently. In this work, we develop a prediction model for video memorability, including complexities of video content in it. Detailed feature analysis reveals that the proposed method correlates well with existing findings on memorability. We also describe a novel experiment of predicting video sub-shot memorability and show that our approach improves over current memorability methods in this task. Experiments on standard datasets demonstrate that the proposed metric can achieve results on par or better than the state-of-the-art methods for video summarization.*

## 1. Introduction

Internet today is inundated with videos. The popular video site, YouTube, alone has more than a billion users and millions of hours of videos being watched every day [1]. Thus, it has become imperative to investigate into advanced technologies for organization and curation of videos. Further, as any such system would involve interaction with humans, it is essential to take cognitive and psychological factors into account for designing an effective system. Moreover, it has been show that metrics like popularity [24] and virality [9] can be predicted by analysing visual features.

An important aspect of human cognition is memorability or the ability to recall visual content after viewing it. Memorability is intricately related to perceptual storage capacity of human memory [3]. Recent studies have further shown

that for prediction of image memorability, deep trained features can achieve near human consistency [25]. There have been also related works in image memorability exploring different features and methods [23, 26, 6]. However, modeling and predicting memorability for video content has not been looked into sufficiently. This is a challenging problem because of added complexities of video like duration, frame rate, etc. Videos also convey multitude of visual concepts to the user, hence, it becomes difficult to ascertain the memorability of the overall content. Further, the temporal structure of the video also needs to be taken into account while modeling video content memorability.

An earlier approach to model video memorability by Han *et al.* [16] deploys a survey-based recall experiment. Here, the participants (about 20) were initially made to watch several videos played together in a sequence, followed by a recall task after two days or a week, where they were asked if they remember the videos being shown. The score for a video was taken to be the fraction of correct responses by the participants. Due to the long time span of experiment, it is difficult to scale it to larger participant size. Further, there is no control over the user behavior between viewing and recall stage. Moreover, the method used fMRI measurements for predicting memorability, which would be difficult to generalize.

To this end, we design an efficient method to compute video memorability, which can be further generalized to applications like video summarization or search. The proposed framework required the participants to complete a survey-based recall task, where they initially watched several videos in a sequence, similar to the earlier approach. However, the recall experiment started after a short rest period of 30s, and the participants were asked textual recall questions, instead of the full video being shown again. The textual questions were constructed from manual annotations of the videos. This was inspired from previous work in human memory research [7, 2], which showed that human

memory stores information semantically. Further, the procedure of a textual questions-based recall survey has been followed in experimental psychology literature [20]. The response time of the user was taken to be the measure of video memorability. Thus, the proposed survey avoids the long gap between viewing and recall stage, hence, making it scalable and efficient, as compared to [16]. We will further release the video memorability dataset in public to help advance research in the field<sup>1</sup>.

We conduct an extensive feature analysis to build a robust memorability predictor. We provide a baseline using state-of-the-art deep learning-based video features. We also explore semantic description, saliency and color descriptors, which have been found to be useful for memorability prediction in prior work on images [11, 25]. Further spatio-temporal features are added to describe video dynamics. We further show that the proposed video memorability model improves over static image memorability in predicting the memorability of video sub-shots (3 – 5s video clips around a video frame). This experiment validates that image memorability is not sufficient to model memorability of video content. We demonstrate application of the model to video summarization task. Videos can be captured for different purposes, with diverse content, duration and quality. Video summarization is, therefore, a challenging task, especially for the content creators who want to ensure that the summary is remembered well by the viewers [29]. In this work, we show that the proposed video memorability framework, which captures human memory recall, can further improve the state-of-the-art in video summarization. The contribution of the work is as follows:

1. We present a novel method for measuring video memorability through a crowd-sourced experiment.
2. We establish memorability as a valuable metric for video summarization and show better or at par performance with the state-of-the-art methods.
3. We demonstrate that proposed image memorability is not sufficient for analyzing memorability of short videos (or sub-shots).
4. We would further release an annotated video memorability dataset, to aid further research in the field.

## 2. Literature Survey

In this section, we discuss the prior work on memorability and related concepts of saliency and interestingness. We also describe the state-of-art work for video descriptors, video semantics and summarization.

**Memorability:** Recent works have explored memorability of images [18, 23, 26, 6, 25, 17]. Memorability of objects in images was studied in [11], while that of natural scenes was explored in [35]. There have been works

<sup>1</sup><https://research.adobe.com/project/video-memorability/>

studying the different aspects of memory, like visual capacity [45, 3] as well as representation of visual scenes [28]. The effect of extrinsic factors on memorability has also been looked into [5]. The recent work by Han *et al.* [16] models video memorability using fMRI measurements.

**Saliency:** Saliency refers to the aspects of visual content which attract human attention. There has been ample work in computing image saliency [19, 46, 21] and its applications to recognition tasks [42]. The saliency feature has been found to be relevant for predicting memorability in [11, 25].

**Interestingness:** Image interestingness was explored by Gygli *et al.* [13], and was extended for video summarization task [15, 14]. Interestingness score of an image was computed as a fraction of users who considered it interesting. However, the interestingness score is subjective and varies considerably with user preferences [13]. Further, the applications to video summarization [15, 14] use varied prediction models for interestingness. On the other hand, we demonstrate that the proposed memorability model can be generalized to different video summarization scenarios. Zen *et al.* [54] described an interestingness model using mouse activities. However, this may not generalize across different viewing conditions (e.g. mobile devices).

**Video descriptors:** Several deep learning-based features have been proposed for video classification task. There have been attempts at extending image-based deep features to videos using different fusion schemes [22, 53]. Tran *et al.* [47] described a 3D-CNN model for action recognition in videos. Fusion of appearance and motion models using deep learning have also been explored [43, 12]. In addition, shallow features like dense spatio-temporal [50] have also been shown to improve classification accuracy when used in conjunction with deep learning features.

**Video semantics:** Inspired by the work on image captioning, there have been recent works in language description of videos. In particular, Donahue *et al.* [10] proposed an LSTM-based approach for video description. This was followed by several works exploiting recurrent network architecture for describing videos [49, 40, 39]. Recently, a method exploiting external information for video captioning was described in [48].

**Video summarization:** There is rich literature on different video summarization techniques. Recently, there has been work exploring sub-modular optimization [15], exemplars [55], object proposals [36] and *Determinantal point processes* (DPPs) [56] for summarization. There have also been related works in summarizing ego-centric videos [34, 31]. A keyword query-based summary method is described in [41].

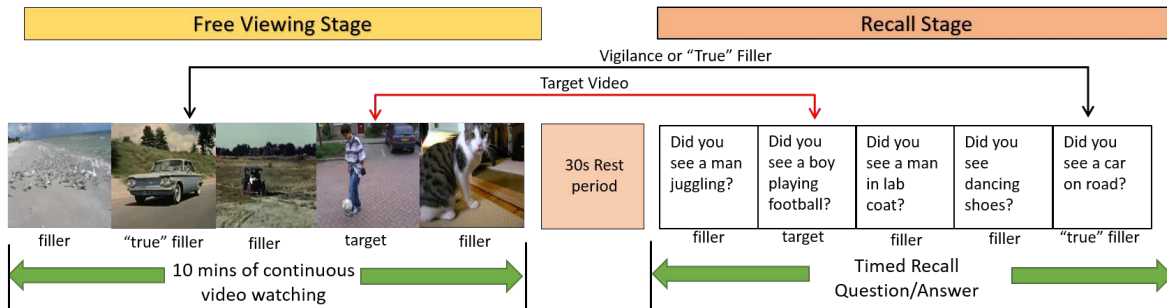


Figure 1: Workflow for the proposed survey design to measure video memorability.

### 3. Video Memorability

In this section, we describe modeling of video memorability, followed by a detailed analysis of memorability results. Figure 1 shows the overall workflow for memorability ground truth collection.

#### 3.1. Ground Truth Collection

The first step to model video memorability is to collect the ground truth of memorability scores. This was done through an Amazon Mechanical Turk (AMT) based crowd-sourced experiment on TRECVID 2012 [38] dataset.

##### 3.1.1 Dataset

TRECVID 2012 [38] consists of about 500 videos taken from the internet archives, ranging across various categories like nature, sports, animal and amateur home videos. The duration of the videos was typically 15 – 30s. For the survey, the videos were first manually captioned to capture the content shown in them. Then, 116 videos were selected across the various categories. These videos were partitioned into two sets - 100 *target* videos and 16 *filler* videos. The target videos were used to construct our model. For the survey, 25 unique combinations of videos, each consisting of 4 different target videos and the same set of 16 filler videos were prepared. Thereafter, 4 permutations of each of these 25 combinations were created, keeping the order and the positions of the filler videos fixed. For the remaining positions, the order of target videos were changed according to the Latin square arrangement [51]. This ensured that each target video was shown at 4 different positions to the users. The overall length of each of these 100 video sequences was about 10 minutes.

##### 3.1.2 Survey Design

We conducted a recall-based experiment on AMT to collect the memorability ground truth. For the experiment, the participants had to first complete watching a full sequence, without browsing away from the survey page. To avoid any

observer effect, the participants were not informed about the recall experiment at the end of free viewing. Further they were not allowed to repeat the survey.

After viewing the video sequence, there was a rest period of 30s and then, the subject was asked 20 yes/no questions. He was given 5s to respond to each of them, and there was no provision of changing the response after the time was over. No response within the 5s duration was treated as a wrong reply. The questions were constructed from the manual text annotation for the video. Some sample questions from the survey are presented in Figure 1. Out of the 20 questions, 8 were true positives, out of which 4 corresponded to the target videos. The rest 4 were randomly chosen from the 16 filler videos, which we call as vigilance or "true" filler videos. The rest of the questions did not relate to any of the shown videos. The questions were randomly ordered for each survey to avoid any systematic bias in response. It was manually ensured that no two textual questions nor any two videos in a sequence were similar in content. The time that the subject took to respond each question was recorded. The survey was conducted with 500 AMT workers, with each sequence permutation being viewed 5 times, hence, giving 20 responses for each target video.

##### 3.1.3 Memorability Score Computation

The memorability score for each video was then calculated as follows:

- First, participants with precision less than 50% were removed from further calculations. This precision was calculated over both target and vigilance videos. This was done to remove users, who may have answered the questions in a random fashion (random behavior precision is about 40% for this setting).
- Consider a target video  $i$  seen by participant  $j$ . Then, memorability score ( $MemScore$ ) of the video  $i$  for par-

ticipant  $j$  is:

$$MemScore(i, j) = \begin{cases} \frac{r(i, j)}{\bar{r}(j)} & \text{correct recall} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where,  $r(i, j)$  is the time left for participant  $j$  in recalling video  $i$ , and  $\bar{r}(j)$  is the mean time left for participant  $j$ , in correctly recalling the videos (including filler videos). For incorrect responses time left was taken to be 0.

- For the video  $i$ , the final memorability score,  $MemScore(i)$ , is the average score across all the participants:

$$MemScore(i) = \frac{\sum_{j=1}^N MemScore(i, j)}{N_i} \quad (2)$$

where,  $N_i$  is number of participants viewing the video.

Note that unlike the hit rate metric in [25, 17], we use a continuous metric based on recall time to capture the strength of memory, inspired from the work of Mickes *et al.* [37]. However, we do find that there is high correlation between hit rate and the proposed metric ( $\rho = 0.91$ ). Further we follow a user-based normalization for score calculation to neutralize background factors like the system used for answering the survey or biases specific to the user, which might affect the user response time.

### 3.2. Memorability Analysis

Here, we analyze the output of the crowd-sourced survey for modeling video memorability.

**Memorability Score Distribution:** Figure 2(a) shows the distribution of video memorability scores across videos. It can be seen that the distribution peaks around 1, while more memorable videos getting scores in the range of 1.3 – 1.5. The overall distribution is skewed, with some videos getting scores as low as 0.5. The high scores around 1.3 – 1.5 means that more memorable videos are recalled faster than the average user recall time. Some of the least and the most memorable videos are shown in Figure 3.

**User Response Time:** Figure 2(b) shows that the distribution of average user recall time has considerable variation. This justifies our choice of user-based normalization for score calculations.

**Effect of video category:** Average memorability scores for different categories of videos are shown in Figure 2(c). It can be seen that *animals* and *objects* videos are most memorable (also as per Figure 3), followed by *human* and *sports* videos. The *nature* and *outdoor* videos have lower scores on an average. Thus, the semantic category of the video also affects its memorability.

**Human Response Consistency:** We also analyzed the consistency of human responses in the AMT survey. The output responses were divided randomly into equal halves, and Spearman’s correlation ( $\rho$ ) was calculated between the memorability score outputs of the two halves. The process was repeated 25 times. We get a high average correlation,  $\rho = 0.68$ , which is consistent with findings in the previous works [17, 11] that memorability is intrinsic to the visual content.

**Effect of complexity of textual questions:** The correlation of Flesch-Kincaid Grade Level readability metric [27] of the textual questions with the memorability scores was found to be quite low ( $\rho = 0.0003$ ). Thus, we don’t observe any effect of complexity of questions on memorability scores.

## 4. Predicting Memorability

In this section, we discuss the task of predicting video memorability. The feature extraction from videos is described in Section 4.1 and an analysis of features for memorability prediction is discussed in Section 4.2.

### 4.1. Feature Extraction

Previous works on memorability [11, 25] have shown semantics, saliency and color to be important features for predicting memorability. Further, we extract spatio-temporal features to represent video dynamics, and provide a baseline using a recent, state-of-art deep learning feature for video classification.

- **Deep Learning (DL):** We extracted the recently proposed C3D deep learning feature [47], trained on the Sports-1M dataset [22] from the videos. The feature has been shown to achieve state-of-the-art classification results on different video datasets. Following the work, we used the activation of the  $fc - \delta$  layer of the pre-trained C3D network to create a 4096-dimensional representation of the video.
- **Video Semantics (SEM):** We used the improved video captioning method developed in Venugopalan *et al.* [48] to first generate the semantic description of the videos. The generated text was then fed to a recursive auto-encoder network [44] to generate a 100-dimensional representation of the videos.
- **Saliency (SAL):** Saliency or the aspect of visual content which grabs human attention has shown to be useful in predicting memorability [17, 11]. We extracted the saliency feature for the video as follows. First, we generated saliency probability maps, using the method proposed in [21], on 10 frames extracted at uniform intervals from the video. This was followed by averaging the saliency maps over the frames, and re-sizing

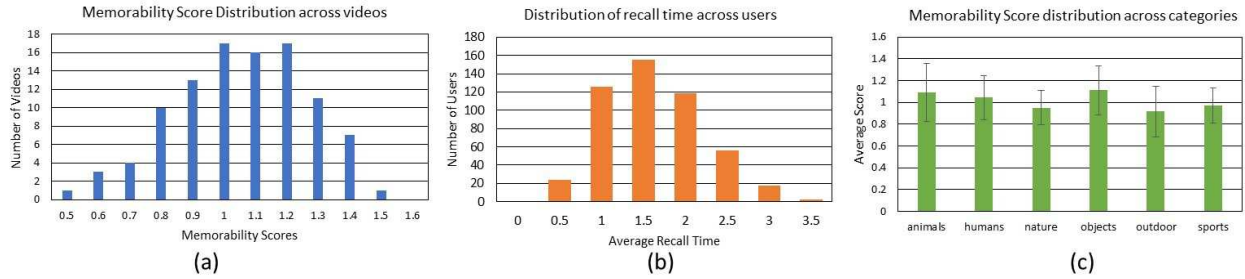


Figure 2: Analysis of the output memorability scores: Distribution of (a) scores across videos, (b) recall time of participants and (c) scores over categories.

the averaged map to  $50 \times 50$ , followed by vectorization, to get the final feature.

- **Spatio-Temporal features (ST):** We used the recent state-of-the-art dense trajectory method [50] to extract a 4000-dimensional vector to represent the spatio-temporal aspect of the video.
- **Color features (COL):** A 100-dimensional color feature was generated for each video by averaging the 50-binned hue and saturation histograms for 10 frames extracted at uniform intervals from the video, followed by concatenation.

## 4.2. Prediction Analysis

Here, we describe the training of regressor for predicting video memorability, and an analysis of importance of different features. For training the regressor, the dataset was randomly split into 80 training videos and the rest 20 for test, and the process was repeated 25 times. We used random forest (RF) regressor to train the model for individual features, tuned using cross-validation. For combining the features, we simply averaged the output regression scores of the individual features. Table 1 reports RMSE for different feature combinations. The results are obtained by averaging over all the 25 runs.

Features	RMSE
COL	$0.155 \pm 0.001$
ST [50]	$0.146 \pm 0.002$
SAL [21]	$0.142 \pm 0.002$
DL (C3D [47])	$0.140 \pm 0.002$
SEM [48]	$0.138 \pm 0.003$
DL+ST+SAL+COL	$0.136 \pm 0.001$
<b>SEM+ST+SAL+COL</b>	<b><math>0.135 \pm 0.001</math></b>

Table 1: Performance analysis of different features. DL: Deep Learning, SEM: Semantics, SAL: Saliency, ST: Spatio-temporal, COL: Color.

**Feature Analysis:** Table 1 shows the performances for different feature combinations. It can be seen that the deep learning-based features, DL and SEM individually achieve low RMSE values, with the latter performing better. Among the shallow features, SAL feature exhibit the lowest RMSE followed by ST and COL features. The better performance of SAL features might be because it captures if the subject of the video grabs human attention or not, as seen in Figure 3. It can be seen that the top 3 memorable videos have salient foreground. However, the *color pencils* video, having similar saliency map as the *tree without leaves* video, has a very different memorability score. Thus, saliency alone is not sufficient to explain memorability. The worse performance for ST might be because video dynamics alone is not sufficient to account for the memorability score. Overall feature combinations further lower the RMSE values.

**Final Memorability Predictor:** The final memorability classifier was trained over all of the 100 target videos, using the SEM+ST+SAL+COL feature combination. We used this regressor for all further experiments.

## 5. Sub-Shot Memorability

In this section, we discuss the problem of predicting sub-shot memorability, and how the existing image memorability work is not sufficient to address the same. We define a sub-shot as a short clip of around 3 – 5s around a selected frame of the video. Due to the short duration, the sub-shot can generally be considered to have homogeneous composition, and that the selected frame is a good representation of the sub-shot. We conducted a survey-based recall experiment to collect the memorability ground truth for sub-shots, following the procedure as described in Section 3.

First, we selected 50 target videos from the TRECVID 2012 [38], different from the target videos used in Section 3. For each video, a sub-shot of 3s around an image frame selected randomly, was extracted. The crowd-sourced AMT survey was designed following the procedure described for

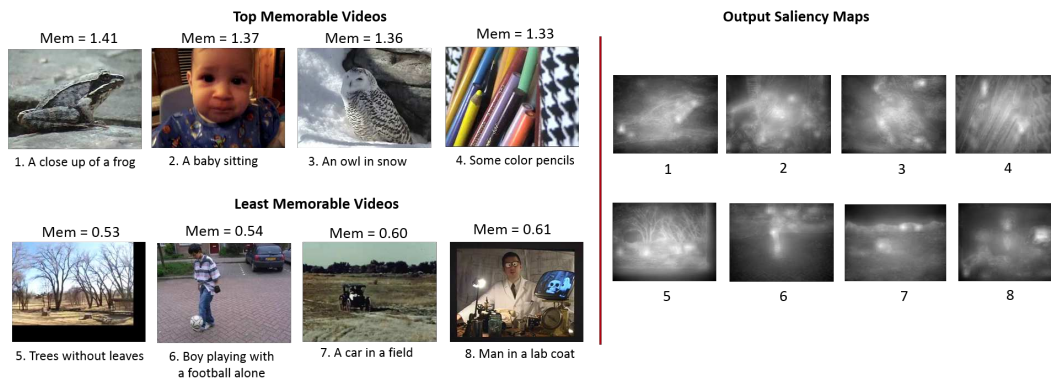


Figure 3: Examples of the most and the least memorable videos from the crowd-sourced video memorability experiment with TRECVID 2012. The output saliency maps [21] for these videos are displayed along with.

video memorability in Section 3, except for a change at the recall stage.

The free viewing sequences consisted of 4 target sub-shots and 16 filler sub-shots, as before. The filler sub-shots were extracted from the filler videos used for the earlier experiment. Due to the shorter video lengths, total free viewing stage lasted for around 2 minutes. During the recall experiment, the participants were asked if they can recall the displayed images (instead of the textual question in original survey). This was done because the sub-shot can be represented by the chosen image frame effectively. For the target as well as “true” filler sub-shots, the corresponding image frame of the video was used, for other slots random images corresponding to none of the shown videos were used. The images were flashed for 0.5s, and then the subject was asked if he can recall the displayed image in 5s. The final score was calculated using the procedure described in Section 3.

**Human Response Consistency:** A consistency analysis of the annotations, similar to the one conducted in Section 3.2 yielded a Spearman’s correlation of 0.45. Thus, sub-shots also have consistent memorability across users, similar to video memorability.

**Prediction Analysis:** We conducted a comparison of the proposed video memorability regressor with the existing image memorability work [25], on predicting the memorability of sub-shots. Image memorability scores were computed by running the pre-trained model from [25] on the selected frame for each sub-shot. Table 2 demonstrates the results of the comparison. It can be seen that image memorability yields much lower Spearman’s correlation value than video memorability regressor. This result demonstrates that complexities of video data must be accounted for, in order to predict memorability. Further, the moderate-to-low correlation for both the cases indicate that further investigations are required into how memorability predictions can be generalized across different kinds of tasks (e.g. video to

sub-shot or image to sub-shot).

Method	Spearman’s cor. ( $\rho$ )
Image Mem. [25]	0.06
Video Mem.	0.20

Table 2: Correlation results for image memorability [25] and the proposed video memorability with the ground truth.

## 6. Video Summarization

In this section, we describe the application of the proposed method to video summarization tasks. Recently a state-of-the-art algorithm for summarization based on supervised learning of sub-modular objective function was proposed by Gygli *et al.* [15]. The framework combined several image-based objectives like interestingness, uniformity and representativeness to improve the quality of video summary. The weights given to each of the objective criteria were learned using a supervised learning algorithm trained using reference human summaries. Here, we further incorporate the proposed video memorability framework as an objective criterion for summarization. We believe this would help improve quality of summaries further.

**Memorability objective:** For a video  $\mathcal{V}$  partitioned into  $N$  segments,  $\{s_i\}_{i=1}^N$ , memorability objective,  $VidMem$  for a selection of  $K \subset \{1, 2, \dots, N\}$  segments is defined as:

$$VidMem = \sum_{i \in K} MemScore(s_i), \quad (3)$$

where,  $MemScore(s_i)$  is the predicted memorability score for segment  $s_i$ . It can be shown that the objective function is sub-modular. Given the functions for scoring summaries on memorability ( $VidMem$ ), uniformity ( $VidUnif$ ) [15] and

representativeness (*VidRep*) [15], the overall objective criteria for selecting the summary,  $\mathbf{y}$  is given as:

$$\mathbf{y}_{opt} = \operatorname{argmax}_{\mathbf{y} \in 2^{\mathcal{V}}, |\mathbf{y}| \leq L} \sum_{f \in \mathcal{F}} w_f f(\mathbf{y}; \mathcal{V}) \quad (4)$$

where,  $L$  is the length of summary,  $\mathcal{F} := \{VidMem, VidRep, VidUnif\}$ ,  $f(\mathbf{y}; \mathcal{V})$  is the summary score using  $f$  and weights  $w$  are learned using the supervised sub-modular optimization as described in [15]. The results are demonstrated on SumMe user [14] and UT Egocentric [30] video datasets.

## 6.1. User Video Dataset

The SumMe user video dataset [14] consists of 25 short videos with lengths ranging from 1 – 7 minutes. The videos depict various activities like sports, cooking, different outdoor activities, traveling, etc. Each video has around 15 (total 390) reference ground truth summaries, generated by humans in a controlled environment. We followed the pre-processing and evaluation protocol described in [15] to have a consistent comparison with the prior art.

**Pre-processing:** The videos were partitioned using super-frame segmentation method [14]. For each segment, SEM, ST, SAL and COL features were extracted, and memorability scores were predicted using the final model (SEM+ST+SAL+COL) trained in Section 4.

**Evaluation:** The dataset was split 12-ways and a leave-one-out method was followed for evaluating the algorithm. The results were averaged over 100 runs. The methods were evaluated for a budget of 15% of the extracted segments. The training was done using the reference user summaries for each video. During the test time, the generated summary was compared with all the reference summaries, and the maximum overlap was taken to get the final F-measure and Recall results, as described in [14].

**Results:** Table 3 shows the comparison of the proposed memorability-based framework with the previous methods. It can be seen that the proposed video memorability alone is able to achieve state-of-the-art F-measure score on the dataset. The results further increase through combination with representativeness and uniformity objectives. Further, the memorability objective gets 96% weight in the supervised training with all the objectives, thus, reinforcing the usefulness of the method. Further, an illustration of summarization achieved by using memorability objective is shown in Figure 4. It can be seen that memorability picks up frames more relevant to users, as well as captures different events in videos well.

## 6.2. UT Egocentric (UTE) Dataset

UTE dataset [30] consists of 4 videos, each with 3 – 5 hours of video content. The video content was recorded through a wearable camera, and logs day activities of the

Methods	F-measure	Recall
UserSum [14]	39.34 ± 0.00%	<b>44.44 ± 0.00%</b>
Uniformity	24.68 ± 0.04%	27.08 ± 0.08%
Representativeness	26.69 ± 0.00%	26.65 ± 0.00%
Interesting [15]	39.52 ± 0.00%	42.50 ± 0.00%
Uni.+Rep.+Int. [15]	39.68 ± 0.09%	43.01 ± 0.08%
Zhang <i>et al.</i> [55]	40.9%	–
Vid. Memorability	41.11 ± 0.10%	37.91 ± 0.11%
Uni.+Rep.+Mem.	<b>41.21 ± 0.12%</b>	38.41 ± 0.15%

Table 3: Evaluation results for summarization with 15% budget on SumMe dataset.

wearer. Thus, the videos may be repetitive and were shot in an uncontrolled fashion. Textual captions for 5s segments of each of the video, as well as 3 reference summaries for each video were provided by Yeung *et al.* [52]. We followed the pre-processing and evaluation protocol described in [15] to have a consistent comparison with the prior art.

**Pre-processing:** The videos were divided into 5s segments and then memorability score was calculated for each segment as described in the previous experiment. For each segment, SEM, ST, SAL and COL features were extracted, and memorability scores were predicted using the final model (SEM+ST+SAL+COL) trained in Section 4.

**Evaluation:** Firstly, for all the videos, segment-based reference summaries were generated using the provided textual summary, following the method proposed in [52]. We used greedy optimization based on bag-of-word model to produce the segment-based reference summaries. The dataset was then split 4-ways and a leave-one-out train/test was followed, similar to [15]. The results were averaged over 100 runs. During the test time, the generated segment-based summary was converted to textual description, and then was compared to the reference text summaries using ROUGE-SU method [33]. The ROUGE-SU computes uni-gram and skip-bigram co-occurrence between candidate and reference summaries, after stemming and removing stop word in the summaries.

**Results:** The proposed method was evaluated for two summary lengths - a shorter length of 1 min 20s and a longer length of 2 mins. The results for evaluations are shown in Table 4 and Table 5. It can be seen that the proposed method achieves the best recall for shorter summary length, whereas for longer summaries, the performance is comparable to Interestingness metric [15]. This may be because we do not employ the manual annotations provided in [30] to identify important objects, which was used in interestingness calculation [15]. Further, with the increase in summary lengths, other metrics like uniformity and representativeness also give results close to memorability. This might be because in typical ego-centric videos, there would be only few “memorable” segments relevant to user. With

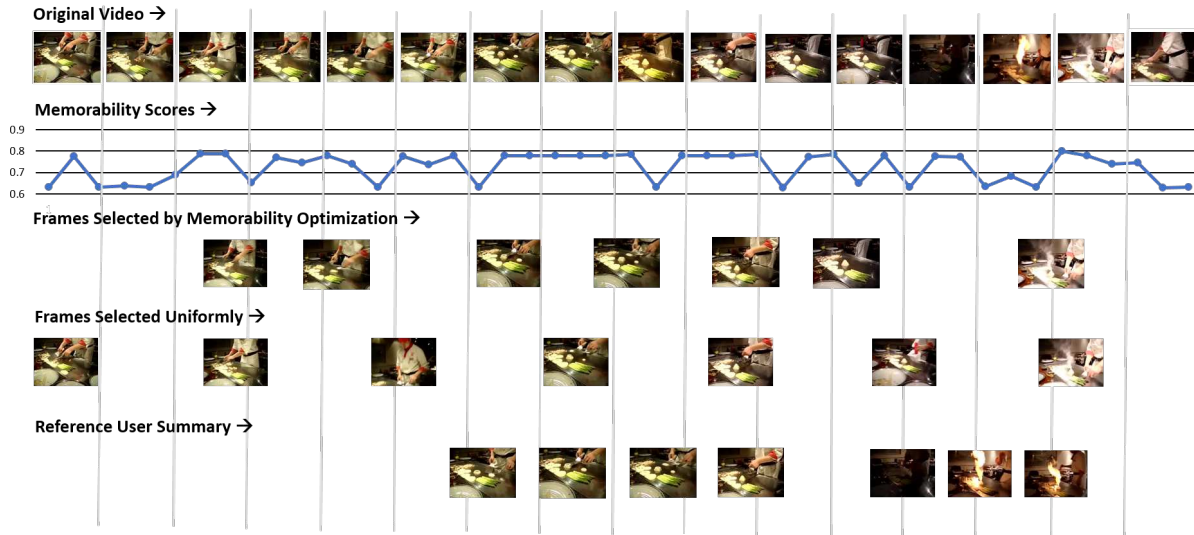


Figure 4: An example of the frame selection through memorability criterion. The shown video from SumMe dataset has a cooking activity going on. As seen in the figure, compared to a uniform selection of frames, memorability criterion picks up frames more relevant to the reference user. It can be seen that memorability score can capture different events and transitions in the video.

the increased budget other metrics can also capture these segments. We believe that the memorability results could be further improved through enhancements in feature design.

Method	F-measure	Recall
Lee <i>et al.</i> [30]	17.40 ± 4.07%	12.20 ± 3.30%
Video MMR [32]	17.73 ± 0.00%	12.49 ± 0.00%
Uniformity	18.75 ± 1.36%	12.92 ± 1.11%
Representativeness	19.08 ± 0.00%	12.95 ± 0.00%
Interesting [15]	20.93 ± 0.00%	15.15 ± 0.00%
Uni.+Rep.+Int. [15]	<b>21.91 ± 0.06%</b>	15.73 ± 0.04%
Vid. Memorability	18.13 ± 0.08%	15.55 ± 0.04%
Uni.+Rep.+Mem.	19.37 ± 0.08%	<b>17.9 ± 0.09%</b>

Table 4: Evaluation results for shorter summarization (1 min 20s) on UTE dataset.

Method	F-measure	Recall
VideoMMR [32]	25.57 ± 0.00%	23.10 ± 0.00%
Uniformity	25.41 ± 1.35%	22.27 ± 1.56%
Representativeness	27.02 ± 0.00%	23.51 ± 0.00%
Interesting [15]	27.07 ± 0.00%	24.78 ± 0.00%
Uni.+Rep.+Int. [15]	<b>29.01 ± 1.18%</b>	<b>26.21 ± 1.23%</b>
Vid. Memorability	26.81 ± 0.04%	21.24 ± 0.02%
Uni.+Rep.+Mem.	28.3 ± 0.5%	23.6 ± 0.3%

Table 5: Evaluation results for longer summarization (2 min) on UTE dataset.

## 7. Conclusions and Future Work

In this work, we have described a robust way to model and compute video memorability. The computed memorability scores are consistent, hence, are intrinsic to the video content, as has been established by prior work in memorability. Further, we analyze different features in predicting memorability, and demonstrate importance of different features. A novel experiment on sub-shot memorability proves that image memorability alone is not sufficient to explain the memorability of sub-shots. Finally, the proposed method achieves state-of-the-art results on different video summarization datasets. This shows that memorability is a viable criteria for creating extractive video summaries. In future, we plan to conduct the video memorability experiment on a larger scale, as well as, design improved features for prediction. This would also require methods proposed in crowd-sourcing literature for addressing ambiguity in the questions and labels [4, 8]. We further believe that application of video memorability to challenging tasks, like video-based recognition or segmentation would enhance the current state-of-the-art.

## 8. Acknowledgement

We thank the anonymous reviewers for their feedback. We also thank our colleagues from Adobe Research who provided insight and expertise that greatly assisted the research. We particularly thank Atanu Ranjan Sinha and P. Anandan for their valuable inputs.



## References

- [1] Statistics - youtube. <https://www.youtube.com/yt/press/statistics.html>. 1
- [2] A. D. Baddeley and G. Hitch. Working memory. *Psychology of learning and motivation*, 8:47–89, 1974. 1
- [3] T. F. Brady, T. Konkle, G. A. Alvarez, and A. Oliva. Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38):14325–14329, 2008. 1, 2
- [4] J. Bragg, D. S. Weld, et al. Crowdsourcing multi-label classification for taxonomy creation. In *First AAAI conference on human computation and crowdsourcing*, 2013. 8
- [5] Z. Bylinskii, P. Isola, C. Bainbridge, A. Torralba, and A. Oliva. Intrinsic and extrinsic effects on image memorability. *Vision research*, 116:165–178, 2015. 2
- [6] B. Celikkale, A. Erdem, and E. Erdem. Visual attention-driven spatial pooling for image memorability. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 976–983, 2013. 1, 2
- [7] A. M. Collins and M. R. Quillian. Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, 8(2):240–247, 1969. 1
- [8] J. Deng, O. Russakovsky, J. Krause, M. S. Bernstein, A. Berg, and L. Fei-Fei. Scalable multi-label annotation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '14*, pages 3099–3102, New York, NY, USA, 2014. ACM. 8
- [9] A. Deza and D. Parikh. Understanding image virality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1818–1826, 2015. 1
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 2625–2634, 2015. 2
- [11] R. Dubey, J. Peterson, A. Khosla, M.-H. Yang, and B. Ghanem. What makes an object memorable? In *IEEE International Conference on Computer Vision*, pages 1089–1097, 2015. 2, 4
- [12] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1933–1941, 2016. 2
- [13] M. Gygli, H. Grabner, H. Riemenschneider, F. Nater, and L. Van Gool. The interestingness of images. In *IEEE International Conference on Computer Vision*, pages 1633–1640, 2013. 2
- [14] M. Gygli, H. Grabner, H. Riemenschneider, and L. Van Gool. Creating summaries from user videos. In *European Conference on Computer Vision*, pages 505–520. Springer, 2014. 2, 7
- [15] 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*, June 2015. 2, 6, 7, 8
- [16] J. Han, C. Chen, L. Shao, X. Hu, J. Han, and T. Liu. Learning computational models of video memorability from fmri brain imaging. *IEEE transactions on Cybernetics*, 45(8):1692–1703, 2015. 1, 2
- [17] P. Isola, J. Xiao, D. Parikh, A. Torralba, and A. Oliva. What makes a photograph memorable? *IEEE transactions on Pattern Analysis and Machine Intelligence*, 36(7):1469–1482, 2014. 2, 4
- [18] P. Isola, J. Xiao, A. Torralba, and A. Oliva. What makes an image memorable? In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 145–152. IEEE, 2011. 2
- [19] L. Itti. Quantifying the contribution of low-level saliency to human eye movements in dynamic scenes. *Visual Cognition*, 12(6):1093–1123, 2005. 2
- [20] L. L. Jacoby, J. P. Toth, and A. P. Yonelinas. Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, 122(2):139, 1993. 2
- [21] T. Judd, K. Ehinger, F. Durand, and A. Torralba. Learning to predict where humans look. In *IEEE International Conference on Computer Vision*, pages 2106–2113. IEEE, 2009. 2, 4, 5, 6
- [22] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. 2, 4
- [23] A. Khosla, W. A. Bainbridge, A. Torralba, and A. Oliva. Modifying the memorability of face photographs. In *IEEE International Conference on Computer Vision*, pages 3200–3207, 2013. 1, 2
- [24] A. Khosla, A. Das Sarma, and R. Hamid. What makes an image popular? In *Proceedings of the 23rd international conference on World wide web*, pages 867–876. ACM, 2014. 1
- [25] A. Khosla, A. S. Raju, A. Torralba, and A. Oliva. Understanding and predicting image memorability at a large scale. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 2, 4, 6
- [26] J. Kim, S. Yoon, and V. Pavlovic. Relative spatial features for image memorability. In *21st ACM international conference on Multimedia*, pages 761–764. ACM, 2013. 1, 2
- [27] J. P. Kincaid, R. P. Fishburne Jr, R. L. Rogers, and B. S. Chissom. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, Naval Technical Training Command Millington TN Research Branch, 1975. 4
- [28] T. Konkle, T. F. Brady, G. A. Alvarez, and A. Oliva. Scene memory is more detailed than you think the role of categories in visual long-term memory. *Psychological Science*, 21(11):1551–1556, 2010. 2
- [29] H. V. Le, S. Clinch, C. Sas, T. Dingler, N. Henze, and N. Davies. Impact of video summary viewing on episodic memory recall: Design guidelines for video summarizations. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4793–4805. ACM, 2016. 2

- [30] 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)*, CVPR '12, pages 1346–1353, 2012. 7, 8
- [31] Y. J. Lee and K. Grauman. Predicting important objects for egocentric video summarization. *International Journal of Computer Vision*, 114(1):38–55, 2015. 2
- [32] Y. Li and B. Merialdo. Multi-video summarization based on video-mm. In *International Workshop on Image Analysis for Multimedia Interactive Services*, pages 1–4. IEEE, 2010. 8
- [33] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*, volume 8. Barcelona, Spain, 2004. 7
- [34] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2714–2721, 2013. 2
- [35] M. Mancas and O. Le Meur. Memorability of natural scenes: The role of attention. In *2013 IEEE International Conference on Image Processing*, pages 196–200. IEEE, 2013. 2
- [36] J. Meng, H. Wang, J. Yuan, and Y.-P. Tan. From keyframes to key objects: Video summarization by representative object proposal selection. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2016. 2
- [37] L. Mickes, J. T. Wixted, and P. E. Wais. A direct test of the unequal-variance signal detection model of recognition memory. *Psychonomic Bulletin & Review*, 14(5):858–865, Oct 2007. 4
- [38] M. M. J. F. G. S. B. S. W. K. A. F. S. P. Over, G. Awad and G. Quénot. Trecvid 2012 – an overview of the goals, tasks, data, evaluation mechanisms and metrics. In *Proceedings of the TRECVID 2012, NIST, USA*, 2012. 3, 5
- [39] P. Pan, Z. Xu, Y. Yang, F. Wu, and Y. Zhuang. Hierarchical recurrent neural encoder for video representation with application to captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1029–1038, 2016. 2
- [40] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4594–4602, 2016. 2
- [41] A. Sharghi, B. Gong, and M. Shah. Query-focused extractive video summarization. In *European Conference on Computer Vision*, pages 3–19. Springer, 2016. 2
- [42] G. Sharma, F. Jurie, and C. Schmid. Discriminative spatial saliency for image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3506–3513. IEEE, 2012. 2
- [43] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014. 2
- [44] R. Socher, J. Pennington, E. H. Huang, A. Y. Ng, and C. D. Manning. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 151–161. Association for Computational Linguistics, 2011. 4
- [45] L. Standing. Learning 10000 pictures. *The Quarterly journal of experimental psychology*, 25(2):207–222, 1973. 2
- [46] A. Toet. Computational versus psychophysical bottom-up image saliency: A comparative evaluation study. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 33(11):2131–2146, 2011. 2
- [47] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *IEEE International Conference on Computer Vision*, pages 4489–4497, 2015. 2, 4, 5
- [48] S. Venugopalan, L. A. Hendricks, R. Mooney, and K. Saenko. Improving lstm-based video description with linguistic knowledge mined from text. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016. 2, 4, 5
- [49] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko. Sequence to sequence-video to text. In *IEEE International Conference on Computer Vision*, pages 4534–4542, 2015. 2
- [50] H. Wang and C. Schmid. Action recognition with improved trajectories. In *IEEE International Conference on Computer Vision*, pages 3551–3558, 2013. 2, 5
- [51] B. J. Winer, D. R. Brown, and K. M. Michels. *Statistical principles in experimental design*, volume 2. McGraw-Hill New York, 1971. 3
- [52] S. Yeung, A. Fathi, and L. Fei-Fei. Videoset: Video summary evaluation through text. *arXiv preprint arXiv:1406.5824*, 2014. 7
- [53] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 4694–4702, 2015. 2
- [54] G. Zen, P. de Juan, Y. Song, and A. Jaimes. Mouse activity as an indicator of interestingness in video. In *International Conference on Multimedia Retrieval*, pages 47–54. ACM, 2016. 2
- [55] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Summary transfer: Exemplar-based subset selection for video summarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1059–1067, 2016. 2, 7
- [56] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman. Video summarization with long short-term memory. In *European Conference on Computer Vision*, pages 766–782. Springer, 2016. 2