Real-time Video Summarization on Commodity Hardware

Wesley Taylor and Faisal Z. Qureshi Faculty of Science, University of Ontario Institute of Technology Oshawa, ON L1G 0C5 Canada {wesley.taylor3|faisal.qureshi}@uoit.net

ABSTRACT

We present a method for creating video summaries in real-time on commodity hardware. Real-time here refers to the fact that the time required for video summarization is less than the duration of the input video. First, low-level features are use to discard undesirable frames. Next, video is divided into segments, and segment-level features are extracted for each segment. Tree-based models trained on widely available video summarization and computational aesthetics datasets are then used to rank individual segments, and top-ranked segments are selected to generate the final video summary. We evaluate the proposed method on The SumMe Video Summarization (SumMe) dataset and show that our method is able to achieve summarization accuracy that is comparable to that of a current state-of-the-art deep learning method, while posting significantly faster run-times. Our method on average is able to generate a video summary in time that is shorter than the duration of the video.

CCS CONCEPTS

• Computing methodologies → Video summarization;

KEYWORDS

Video summarization, video analysis

ACM Reference Format:

Wesley Taylor and Faisal Z. Qureshi. 2018. Real-time Video Summarization on Commodity Hardware. In *International Conference on Distributed Smart Cameras (ICDSC '18), September 3–4, 2018, Eindhoven, Netherlands.* ACM, New York, NY, USA, Article 4, 9 pages. https://doi.org/10.1145/3243394. 3243689

1 INTRODUCTION

Cameras are now ubiquitous. This has resulted in an explosive growth in user-generated images and videos. In the case of videos, at least, our ability to record videos has far outpaced methods and tools to manage these videos. A skier, for example, can easily record many hours of video footage using an action camera, such as a Go-Pro. Raw video footage, in general, is *unviewable*—the recorded video needs to be summarized or edited in some manner before it can be shared with others. Clearly, no one is interested in watching many hours of skiing video when most of it is bound to be highly repetitive. Manual video editing and summarizing is painstakingly

ICDSC '18, September 3-4, 2018, Eindhoven, Netherlands

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6511-6/18/09...\$15.00

https://doi.org/10.1145/3243394.3243689

slow and tedious. Consequently a large fraction of recorded footage is never shared or even viewed. We desperately need one-touch video editing tools capable of generating video summarizes that capture the meaningful and interesting portions of the video, discarding sections that are boring, repetitive or poorly recorded. Such tools will revolutionize how we share video stories with friends and family via social media.

A meaningful video summarization needs to take into account both the user context and the video content. Two different users may find entirely different sections of a recorded video interesting. Consider, for example, the scenario where someone records a children soccer match. Parents may only be interested in a section in video that shows their child. We refer to this as user context. Video summarization algorithms, therefore, should take into account the likes and dislikes of the viewers of the video summary. Video content is also important. By necessity video summarization algorithms relies upon video content to select which portions of the videos *make the cut*.

This paper develops a real-time video summarization system (Figure 1). The proposed system is able to perform video summarization at speeds that far exceed those achieved by state-of-the-art deep learning approaches for video summarization. We list these approaches in the next section. The proposed system exploits lowlevel image features to efficiently discard segments with *low interestingness* or having *poor quality*. This means that subsequent summarization steps, which are computationally expensive, only deal with the remaining segments. This can lead to significant savings, especially for long duration videos, such as the all day ski trip video in the example mentioned above. A key feature of the proposed system is its ability to generate alternate summaries almost instantaneously. A user can guide the system to generate a different summary thereby injecting user-preference into the process of summarization. Figure 1 shows our summarization pipeline.

We evaluate the proposed method on SumMe video summarization benchmark, and compare our method with a number of existing video summarization schemes. Our method achieves the



Figure 1: Our summarization pipeline. Video V is processed to construct a set of segments S_V . Next, a 124-dimensional feature vector X_s^{124} is extracted for each segment $s \in S_V$. These features are processed to assign a ranking to each segment $s \in S_V$. Final step consists of selecting the top ranked segments to create the summary U_V .

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

highest F_1 -measure. It also achieves highest accuracy on over 50% of the tested videos. We also show that the summarization times of the proposed method increases linearly with the duration of the input video.

The rest of the paper is organized as follows. We briefly discuss related work in the next section. Section 3 discusses video segmentation. Segment ranking is covered in Section 4. The following section describes video summarization. We conclude the paper with evaluation and results and conclusions in the last two sections.

2 BACKGROUND

A majority of the existing video summarization methods follow a common recipe: step 1) video segmentation, step 2) segment ranking and step 3) segment selection [3, 6, 12, 21]. Methods vary in how segmentation is performed and how individual segments are ranked. [24] is an exception to this rule that uses recent advances in deep learning and provides an end-to-end system for video summarization. This method relies upon the availability of suitable training data. Early video summarization methods were unsupervised [3, 12, 21]; however, with the recent availability of high-quality video summarization datasets, many newer methods are supervised [6, 24].

Video summarization has also been explored in the context of robotics [5]. Their motivation stems from the fact that transmitting raw video footage, say to a base station, incurs large communication costs. It is also infeasible in situations where bandwidth is limited. They leverage topic modeling to identify the novel segments of the recorded video with a view to construct a video summarization that captures the salient pieces of the video.

Clustering [12] and attention [3] methods are often used as baselines when evaluating new summarization methods. The first method performs clustering to get segmentation, and uses a 0/1 knapsack for segment selection for final summary generation. The second method extracts attention features for each frame, assigning an interestingness score to each frame. Frames with high interestingness scores are selected to generate the summary. We refer the kind reader to the respective publications for technical details. Suffice to say that both classes of methods are unsupervised and are able to achieve higher accuracy when compared to a method that picks frames (or segments) at random when generating a video summarization. Recent methods outperform both these methods.



Figure 2: Using luminance (0.02), sharpness (3.42) and uniformity (0.15) to label undesirable frames. The values shown in red indicates that these fall below the empirically selected threshold values.

Method developed in [6] is of particular interest to us. [6] not only developed a new method for video summarization. It also created a first-of-its-kind benchmark for video summarization. This dataset is referred to as the SumMe dataset. We too use this dataset to evaluate the performance of our method. [6] uses change point detection for segmentation. These segments are subsequently ranked and the final summary is generated using a 0/1 knapsack formulation. [21] method is similar to the method proposed in [6]. The key difference is that [21] method uses a different set of features for ranking segments.

The current best performing video summarization method is [24]. It uses convolutional and recurrant layers that operate upon sequences of frames and compute interestingness score for each frame. Specifically, this method uses pool-5 layer of GoogLeNet model as frame-level features, which are fed into LSTM units to generate frame and segment level interestingness scores. The key idea is to capture temporal relationship between successive frames to compute frame-level interestingness score suitable for video summarization.

3 VIDEO SEGMENTATION

The algorithm begins by identifying frames that are too dark, blurry, or uniform (see Figure 2). Luminance (Y), sharpness (S) and uniformity (U) values are computed for each frame to label the frame accordingly. Luminance is given by

$$Y = \text{mean}(0.2126 \cdot R + 0.7152 \cdot G + 0.0722 \cdot B)$$

sharpness is computed as

$$S = \text{mean}(G_x^2 + G_u^2)$$
, and

uniformity value is computed by first constucting a normalized 1D grayscale histgoram H with 128 bins and then computing the ratio between the top 5th percentile bins of H and the rest of H. These features have low computational overhead. The algorithm thus avoids wasting precious computational resources (during the subsequent steps) on frames that will not make the final cut any ways.

Next, input video V is divided into one or more non-overlapping segments

$$S_V = \{s_0, \ldots, s_k\}.$$

While these segments do not overlap, we allow for gaps between adjacent segments, i.e., we only require that $end(s_i) < start(s_{i+1})$. We formulate our video as a multidimensional time-series, allowing us to cast video segmentation as a multiple change point detection problem [2, 22].

Change point detection operates upon a time series feature matrix X, where column *i* stores features extracted from frame *i*. Our method extracts 2200-dimensional feature vector from each video frame. Specifically each frame is represented using a HSV histogram with 128 bins per channel and an edge orientation and magnitude histogram with 30 bins each. These features are extracted over a two-level pyramid consisting of 5 regions, which yields a 2200-dimensional feature. Each video is now represented as a $2200 \times n$ matrix X. Here *n* indicates the number of frames. A set of sparse coefficients $\mathbf{A} \in \mathbb{R}^{n \times n}$ is computed from X by solving the following

convex optimization problem.

$$\underset{\mathbf{A}}{\arg\min} \|\mathbf{X} - \mathbf{X}\mathbf{A}\|_{\mathrm{F}}^{2} + \frac{\lambda}{2} \|\mathbf{A}\|_{2,1}.$$
 (1)

A is used to assign a score to each frame, and the top ranked kframes are selected as split points to generate k + 1 segments. We set the problem so that average segment duration is roughly 5 seconds. This method can be thought of as a more robust version of threshold-based and content-aware sampling. Rather than relying simply on local color or brightness features, a combination of color and edge histograms are used to locate segment boundaries based on the statistical properties of the entire video.

Next we refine the segmentation by removing dark, blurry or uniform frames. Segments having a large fraction of undesirable frames are discarded in the process, which also results in further savings down the line. Segments can also be trimmed, discarding undesirable frames at the either end, or split into two or more segments. Adjacent segments containing too few frames are also merged to form a single segment at this stage. This process is shown in Algorithm 1.

4 SEGMENT RANKING

Once all candidate segments S_V for our video V have been located, the next step is to rank these segments. The algorithm begins by extracting frame-wise features, which are subsequently used to rank the individual segments.

4.1 Frame-Level Features

We compute a 62 dimensional feature vector X_{ℓ}^{62} for each frame as follows. The first 59 dimensions correspond to computational aesthetic features computed at each frame (Table 1). We refer the interested reader to [13, 18, 19, 25] for technical details about these features. Dimensions 60 contains the number of faces seen in this frame, and dimension 61 records the number of "salient" faces seen in this frame. The last dimension stores a 1 if the frame is deemed aesthetically pleasing (see below).

4.1.1 Salient Face Detection. The process of finding salient faces consists of three steps: a) face detection, b) (face) feature vector extraction, and c) (face) clustering. The algorithm employs Felzenszwalb's HOG (FHOG) for face detection [4]. To extract a face feature vector, we employ a modfied version of ResNet-34 [8], containing only 29 layers and half the number of filters in each layer. We train the network using a metric loss function over 3 million faces from the FaceScrub [15] and VGG-Face [17] datasets. This model is able to predict with 99.38% accuracy if two faces belong to the same individual on the Labeled Faces in the Wild (LFW) [11] dataset.

Face feature vectors are clustered using Chinese whispers graph clustering algorithm [1]. Chinese whispers is a linear-time hard partitioning, randomized, flat clustering method. A linear-time algorithm is highly desirable since an hour long video can easily contain more than 50,000 face feature vectors. Clustering ensures that each "person" ends up in at most one cluster. Clusters with large memberships identify salient persons. Note that this method requires no prior knowledge about salient faces.

Taylor and Qureshi

Algorithm 1 The algorithm	used for performing	segment merging
and elimination.		

Inputs:

S: A segmentation consisting of *n* segments $\{s_0, \ldots, s_{n-1}\}$

 $d_{\rm m}$: The minimum segment frame duration threshold

on threshold

 $d_{\rm m}, d_{\rm h}$)

	8
	d_{b} : The between segment frame duration
1:	function PostProcessShortSegments(S,
2:	for s_p , s , s_n in $ZIP(S, S[1 :], S[2 :])$ do
3:	if $frames(s) > d_m$ then
4:	continue
5:	end if
6:	$merged \leftarrow False$
7:	if $distance(s_p, s) \le d_b$ then
8:	$S \leftarrow \text{Remove}(S, s_p)$
9:	$start(s) \leftarrow start(s_p)$
10:	merged \leftarrow True
11:	end if
12:	if $distance(s, s_n) \le d_b$ then
13:	$S \leftarrow \text{Remove}(S, s_n)$
14:	$end(s) \leftarrow end(s_n)$
15:	merged \leftarrow True
16:	end if
17:	if merged = FALSE then
18:	$S \leftarrow \text{Remove}(S, s)$
19:	end if
20:	end for
21:	return S
22:	end function

Output: A new version of S with segment merging and elimination applied

The initial "graph" used as input to the clustering algorithm is constructed by simply looping over every pair of features

$$\left| \left(X_{f_a}^{128}, X_{f_b}^{128} \right) \middle| f_a, f_b \in \textit{frames}(s), f_a \neq f_b \right|$$

across all segments and frames computed in the previous step, and creating an "edge" between two nodes when their distance is below some threshold value τ . A value of $\tau = 0.6$ was selected, as it matches the value that was used for the metric loss layer of the deep neural network used in the previous step.

4.1.2 Aesthetic Score. The last dimension contains an aesthetic score of 0 or 1 for this frame. We use an XGBoost classifier trained on A Large-Scale Database for Aesthetic Visual Analysis (AVA) [14] dataset to compute this score. Each image in AVA dataset has an associated user score between 0 and 1, which captures the aesthetic appeal of that image. For our purposes, we assign a score of 0 for any image with ranking less than 0.5. Images with ranking more than 0.5 are assigned a score of 1. We train an XBGoost classifier using 10-fold cross-validation and a train/test split of 70%/30%. The input to this classifier are computational aesthetic features listed in Table 1. The XGBoost classifier obtains an accuracy of 73.66%, which is significantly higher than the reference model shown in [14]. The accuracy of reference model is 53.85%. ILGnet [10] posts the current best accuracy of 82.66%. ILGnet is a deep learning based model, which is more tricky to train and has significantly worse runtime performance than our XGBoost classifier.

4.2 Segment Features

The proposed method computes segment-level features by aggregating frame-level features extracted from frames belonging to each segment. Recall that each frame f is represented as a 62 dimensional feature X_f^{62} . Segment-level feature for each segment $s \in U_V$ is

$$X_s^{124} = \bigcup_{i=0}^{61} \left\{ \operatorname{mean}\left(\left\{ X_f^i \mid f \in s \right\} \right), \operatorname{std}\left(\left\{ X_f^i \mid f \in s \right\} \right) \right\}.$$
(2)

4.3 Ranking

We studied three models—(1) decision trees, (2) random forests, and (3) XGBoost—for ranking segments using the segment features discussed in the previous section. We trained interestingness prediction models for each of the above using segment-level features extracted from videos available in SumMe and Summarizing Web Videos using Titles (TVSum50) datasets. For training purposes these videos are divided into 5 second segments, and segmentlevel features are extracted for each segment. Train-test splits are generated using 10-fold cross-validation on shuffled data, and the mean-squared-error is used as the error metric for evaluating each model. The results for each model are presented in Table 2.

As we can see from Table 2, both the XGBoost and random forest models obtain very similar error rates, with XGBoost slightly out-performing the random forest model, and both significantly out-performing the decision tree model. For this reason, we will use both XGBoost and random forest models for evaluating our system.

4.4 Feature Importance

It is straightforward to compute feature importance when using Decision Trees and XGBoost. In order to see the efficacy of our choice of features, we performed feature importance analysis. Feature importance values are normalized between 0 and 1. A value of 1 suggests that this feature plays an important role within the model. Similarly, a value of 0 indicates that this feature is rarely used during the prediction task. Figure 3 plots feature importance for XGBoost model.

One important conclusion we can draw from Figure 3 is that among all the features used by our model, face detection and recognition features have the least average importance. These features, incidently, are computationally expensive to compute. Our initial hypothesis was that the computational cost of these features would be offset by their actual importance when computing a segment ranking. Figure 3 shows that this is obviously not the case. We,



Figure 3: A plot of feature importances for each feature included in our final feature vector. For the purpose of visualization, we have grouped the features into four major groups, each represented by its own color; blue represents the mean values of aesthetic features, green the variances of theese aesthetic features, red the mean and variance of our XGBoost aesthetics model values, and finally purple the mean and variance values for our face detection and face recognition features. The background of each group additionally contains an aggregate bar which shows the average importance across the entire group.

therefore, decided to exclude face detection and recognition features during segment ranking. This leaves a 120 dimensional feature for segment ranking: $\{X_s^{120} \mid s \in S_V\}$.

Figure 3 suggests that features constructed using XGBoost predictions have the highest average importance score. Recall that XGBoost model is trained on AVA dataset. This means that we are able to train a supervised model for individual image aesthetics and successfully apply this model to the task of segment ranking within the context of video summarization.

5 VIDEO SUMMARIZATION

The final summary U_V leverages segment rankings $\{R_s \mid s \in S_V\}$ computed previously. We formulate segment selection as a 0/1 knapsack problem. Given a set of items (segments) $s \in S_V$, each with a weight (duration) *frames*(*s*) and a value (ranking) R_s , we determine which segments to include in our final summary such that the final length is less than or equal to our target summary duration, and the sum of segment rankings is maximized. Mathematically, we can describe this as

$$\underset{U \subseteq S_V}{\arg\max} \sum_{s \in U} R_s \quad \text{subject to:} \quad \sum_{s \in U} \textit{frames}(s) \leq W.$$

This can be solved via dynamic programming [23]. Define *T* as an $n \times W$ array, and T[i, w] as the maximum score that can be obtained with duration up to or less than *w* using the first *i* items of $S_V = \{s_0, \ldots, s_{n-1}\}$. We get the following recursive definition:

$$T(0,w)=0$$

$$T(i, w) = \begin{cases} T[i - 1, w] \\ \text{if } frames(s_i) > w \\ \max(T[i - 1, w], T[i - 1, w - frames(s_i)] + R_{s_i}) \\ \text{if } frames(s_i) \le w. \end{cases}$$

The solution can be found by computing the value of T[n, W].

ICDSC '18, September 3-4, 2018, Eindhoven, Netherlands

Taylor and Qureshi

Feature	Dim.	Description
Contrast	1	The ratio between the luminance range and average luminance.
Image Mean HSV	3	The average H, S, and V values over the entire image.
Center Mean HSV	3	The average H, S, and V values for the image center quadrant.
Itten Histograms [13]	20	Histograms of H values over 12 bins, S values over 5 bins, and V values over 3 bins.
Itten Contrasts [13]	3	Standard deviation of each Itten Histogram.
Pleasure, Arousal, Dominance [13]	3	Approximate emotional values computed as linear combinations of the mean V and S values.
Haralick Texture Features [7]	13	Average Haralick texture features over all four directions.
Contrast Balance	1	Distance between the original and contrast-normalized grayscale image.
Exposure Quality	1	Negative absolute value of luminance histogram skew.
JPEG Quality [20]	1	No-reference quality estimation algorithm for JPEG images.
Tenengrad [16]	1	Sharpness according to the Tenengrad method.
Spectral Residual [9]	9	Rule of thirds using spectral saliency in 9 quadrants.

Table 1: Low-level aesthetic features extracted from each frame.

Model	Min	Max	Mean	Std. Dev.
Decision Tree	0.04005	0.05145	0.04559	0.00380
Random Forest	0.02302	0.03025	0.02673	0.00238
XGBoost	0.02244	0.02907	0.02537	0.00214

Table 2: Mean-squared-error of each of our three base models evaluated using 10-fold cross validation. We can see that of the three models, XGBoost has the best performance, with the random forest model performing slightly worse, and the decision tree significantly worse.

6 EVALUATION AND RESULTS

We evaluate the proposed method using pairwise F₁-measure on SumMe dataset. SumMe contains multiple summaries from different users, and we need a mechanism for comparing the summary generated by our method with these user-generated summaries. [6] proposed pairwise F₁-measure to perform this comparison and evaluate the performance of a summarization scheme. F₁-measure is computed as follows. Given a summary U and a set of a set of user-generated summaries $J = \{U^0, \ldots, U^n\}$, for each U^i in J compute

$$p_i = \frac{|\text{frames}(U) \cap \text{frames}(U^i)|}{|\text{frames}(U^i)|}$$

and

$$r_i = \frac{| frames(U) \cap frames(U^i) |}{| frames(U) |}$$

Pairwise F₁-measure is then

$$F_U = \frac{1}{n+1} \sum_{i=0}^n 2 \cdot \frac{p_i r_i}{p_i + r_i}$$

For the Random Forest and XGBoost models from Section 4.3, we perform grid search over various model parameters, and continue with the optimal parameters for each variable. In the end, we compare the final pairwise F_1 -measure measures between the Random Forest and XGBoost models, and select the model which attains the highest value. Better methods are represented by higher F_1 -measure values.

Using the default parameters for our XGBoost model, our method obtains an average F_1 -measure value of 0.198. Average F_1 -measure scores obtained by competing methods in [6] and [22] on SumMe dataset are 0.234 and 0.2655, respectively. We fine-tuned the XG-Boost model for segment ranking. The following parameters were considered during grid search: max depth, minimum child weight, gamma, subsample and col-sample by-tree. For our dataset, the optimal values for max depth, minimum child weight, gamma, subsample and col-sample by-tree are 3, 5, 0, 1 and 1, respectively. F_1 -measure was improved from 0.198 to 0.237 using these values.

6.1 Accuracy on SumMe Dataset

We now compare our model to existing techniques on SumMe dataset. Table 3 lists accuracy values for various methods on SumMe dataset. Our method achieves the highest average F_1 -measure among the 5 computational video summarization schemes listed here. Average F_1 -measure scores are provided for different videos in the SumMe dataset. Our method posts the highest scores for roughly 50% of the tested videos.

6.2 Performance

We performed video summarization for each video in the SumMe dataset using our method and recorded the times needed to generate the summaries. These times are shown in Table 4. Notice that Real-time Video Summarization on Commodity Hardware

	Ľ	Humans			Computational Methods					
Videoname	Random	Upper Bound	Worst	Mean	Best	Uniform	Cluster	Attn.	Summe	Ours
Air Force One	0.144	0.490	0.185	0.332	0.457	0.161	0.143	0.215	0.318	0.362
Base jumping	0.144	0.398	0.113	0.257	0.396	0.168	0.109	0.194	0.121	0.106
Bearpark climbing	0.147	0.330	0.129	0.208	0.267	0.152	0.158	0.227	0.118	0.261
Bike Polo	0.134	0.503	0.190	0.322	0.436	0.058	0.130	0.076	0.356	0.301
Bus in Rock Tunnel	0.135	0.359	0.126	0.198	0.270	0.124	0.102	0.112	0.135	0.147
Car railcrossing	0.140	0.515	0.245	0.357	0.454	0.146	0.146	0.064	0.362	0.192
Cockpit Landing	0.136	0.443	0.110	0.279	0.366	0.129	0.156	0.116	0.172	0.201
Cooking	0.145	0.528	0.273	0.379	0.496	0.171	0.139	0.118	0.321	0.348
Eiffel Tower	0.130	0.467	0.233	0.312	0.426	0.166	0.179	0.136	0.295	0.088
Excavators river crossing	0.144	0.411	0.108	0.303	0.397	0.131	0.163	0.041	0.189	0.231
Fire Domino	0.145	0.514	0.170	0.394	0.517	0.233	0.349	0.252	0.130	0.169
Jumps	0.149	0.611	0.214	0.483	0.569	0.052	0.298	0.243	0.427	0.542
Kids playing in leaves	0.139	0.394	0.141	0.289	0.416	0.209	0.165	0.084	0.089	0.093
Notre Dame	0.137	0.360	0.179	0.231	0.287	0.124	0.141	0.138	0.235	0.107
Paintball	0.127	0.550	0.145	0.399	0.503	0.109	0.198	0.281	0.320	0.213
Playing on water slide	0.134	0.340	0.139	0.195	0.284	0.186	0.141	0.124	0.200	0.218
Saving dolphines	0.144	0.313	0.095	0.188	0.242	0.165	0.214	0.154	0.145	0.128
Scuba	0.138	0.387	0.109	0.217	0.302	0.162	0.135	0.200	0.184	0.140
St Maarten Landing	0.143	0.624	0.365	0.496	0.606	0.092	0.096	0.419	0.313	0.557
Statue of Liberty	0.122	0.332	0.096	0.184	0.280	0.143	0.125	0.083	0.192	0.259
Uncut Evening Flight	0.131	0.506	0.206	0.350	0.421	0.122	0.098	0.299	0.271	0.081
Valparaiso Downhill	0.142	0.427	0.148	0.272	0.400	0.154	0.154	0.231	0.242	0.288
car over camera	0.134	0.490	0.214	0.346	0.418	0.099	0.296	0.201	0.372	0.408
paluma jump	0.139	0.662	0.346	0.509	0.642	0.132	0.072	0.028	0.181	0.334
playing ball	0.145	0.403	0.190	0.271	0.364	0.179	0.176	0.140	0.174	0.151
Average	0.139	0.454	0.179	0.311	0.409	0.143	0.163	0.167	0.234	0.237

Table 3: F₁-measure values resulting from testing various summarization methods on videos from SumMe dataset. For each video, among the computational methods, the three highest results are highlighted using different shades of green. Darker shades are used for higher F₁-measure values, and hence better results.

summarization times are smaller than the duration of the videos. The third column shows the speed of video summarization process. On average our method achieves a speed of 1.82 times the actual duration of the video. In other words the time it takes to summarize a video is on average 0.55 times the duration of the video. Figure 4 plots summarization times vs. video durations. It suggests a linear relationship between summarization times and video durations. We fit a first-degree polynomial to this data. The coefficient of determination for this fit is $R^2 = 0.943$, suggesting that a line is a good estimator for this data.

Figure 5 plots average performance vs. accuracy for different methods. A performance value of 1.0 indicates that the summerization time is the same as the duration of the video. We desire methods with performance greater than 1.0. We can view these methods as *faster than real-time*. Newer, computationally expensive methods—SumMe and LSTM—achieve high summarization accuracy; however, these methods posts poor performance. Older, simpler methods on the other hand show high performance scores. These methods, however, have low accuracy scores. Our method is able to achieve high scores for both performance and accuracy.



Figure 4: A plot of the video duration versus computation time data from Table 4. We additionally plot a line of best fit to our data, demonstrating the fact that the complexity of our method appears to be linear in terms of the duration of a video.

Only the LSTM method is able to achieve a higher accuracy score than our method; however, the LSTM method has significantly ICDSC '18, September 3-4, 2018, Eindhoven, Netherlands

Video Name	Duration (s)	Time (s)	Speed
Jumps	38.00	19.12	1.99x
Cooking	85.80	22.16	3.87x
Fire Domino	53.73	27.99	1.92x
St Maarten Landing	70.04	36.72	1.91x
Scuba	74.03	48.45	1.53x
paluma jump	85.89	46.89	1.83x
Bike Polo	102.13	69.50	1.47x
Playing on water slide	102.27	54.76	1.87x
playing ball	103.97	54.52	1.91x
Kids playing in leaves	106.34	71.29	1.49x
Bearpark climbing	133.64	78.31	1.71x
Statue of Liberty	154.52	69.89	2.21x
car over camera	146.21	71.04	2.06x
Air Force One	179.76	103.59	1.74x
Notre Dame	192.00	106.87	1.80x
Base jumping	157.79	105.27	1.50x
Eiffel Tower	198.84	118.90	1.67x
Car railcrossing	169.34	115.14	1.47x
Bus in Rock Tunnel	171.10	109.00	1.57x
Valparaiso Downhill	172.77	115.51	1.50x
Paintball	254.25	137.37	1.85x
Saving dolphines	222.99	120.15	1.86x
Cockpit Landing	301.83	200.50	1.51x
Uncut Evening Flight	322.72	215.42	1.50x
Excavators river crossing	388.84	210.87	1.84x
Average			1.82x

Table 4: Raw performance data for our method applied to each video in the SumMe dataset. The duration of each video is provided, along with the time required for our method to complete, and corresponding speed as a multiplier of the duration of the video.

lower average performance than our method. Figure 6 shows summarization results for our method on a selection of videos taken from the SumMe dataset.

7 CONCLUSIONS

We propose a high performance video summarization system which is able to perform video summarization in an online fashion on commodity hardware. The results demonstrate that our method is able to acquire comparable summarization quality at a fraction of a computational costs of a state-of-the-art LSTM method. Our method, for example, is able to create video summaries of arbitrary duration on a commodity desktop—a i5-3380M CPU and with 16GB of RAM and no dedicated GPU—at times less than the duration of the videos. This suggests that our method may be ideally suited for mobile deployment.

The primary limitation of our method stems from how features are computed for each segment. We have chosen low-level features, which are computationally inexpensive to extract. A downside is that these features are fundamentally limited in terms of capturing semantic information present in a video. We aim to solve this



 $\begin{array}{c} 0.24 \\ \hline 0.22 \\ \hline 0.22 \\ \hline 0.16 \\ 0.14 \\ \hline 0 \\ 0.5 \\ 1 \\ 1.5 \\ 2 \\ 2.5 \\ \hline \end{array}$

Figure 5: Average performance vs. accuracy. Performance is the ratio of the video duration and summarization time. A performance score of greater than 1.0 suggests that summerization times are less than video duration, i.e., it takes less time to summarize a video than it is to record this video. Higher performance valus are highly desireable. Accuracy scores are average F₁-measure. This plot also include performance and accuracy scores of a state-of-the-art LSTM-based method [24].

shortcoming in the future by incorporating additional features into our framework. We are also investigating methods to adapt our framework to incorporate user preferences when creating video summaries.

REFERENCES

Uniform

- Chris Biemann. 2006. Chinese whispers. In Proceedings of TextGraphs: the First Workshop on Graph Based Methods for Natural Language Processing on the First Workshop on Graph Based Methods for Natural Language Processing - TextGraphs '06. Association for Computational Linguistics. https://doi.org/10.3115/1654758. 1654774
- [2] Kevin Bleakley and Jean-Philippe Vert. 2011. The group fused lasso for multiple change-point detection. arXiv preprint arXiv:1106.4199 (2011).
- [3] Naveed Ejaz, Irfan Mehmood, and Sung Wook Baik. 2013. Efficient visual attention based framework for extracting key frames from videos. Sig. Proc.: Image Comm. 28, 1 (2013), 34–44. https://doi.org/10.1016/j.image.2012.10.002
- [4] David A. Forsyth. 2014. Object Detection with Discriminatively Trained Part-Based Models. *IEEE Computer* 47, 2 (2014), 6–7. https://doi.org/10.1109/MC.2014. 42
- [5] Yogesh Girdhar. 2014. Unsupervised Semantic Perception, Summarization, and Autonomous Exploration for Robots in Unstructured Environments. Ph.D. Dissertation. McGill University.
- [6] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc J. Van Gool. 2014. Creating Summaries from User Videos. In Computer Vision - ECCV 2014 -13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VII (Lecture Notes in Computer Science), David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.), Vol. 8695. Springer, 505–520. https://doi. org/10.1007/978-3-319-10584-0_33
- [7] Robert M. Haralick, K. Sam Shanmugam, and Its'hak Dinstein. 1973. Textural Features for Image Classification. *IEEE Trans. Systems, Man, and Cybernetics* 3, 6 (1973), 610–621. https://doi.org/10.1109/TSMC.1973.4309314
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. IEEE Computer Society, 770–778. https://doi.org/10.1109/CVPR.2016.90

Real-time Video Summarization on Commodity Hardware

ICDSC '18, September 3-4, 2018, Eindhoven, Netherlands

Uncut Evening Flight



Figure 6: Our method on a subset of SumMe dataset. For each video, the top plot shows the segment ranking computed by the proposed method. The middle plot shows the ground truth frame-level user rankings. The last plot shows a selection of frames in the summary generated by the proposed method.

- [9] Xiaodi Hou and Liqing Zhang. 2007. Saliency Detection: A Spectral Residual Approach. In 2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007), 18-23 June 2007, Minneapolis, Minnesota, USA. IEEE Computer Society. https://doi.org/10.1109/CVPR.2007.383267
- [10] Xin Jin, Jingying Chi, Siwei Peng, Yulu Tian, Chaochen Ye, and Xiaodong Li. 2016. Deep image aesthetics classification using inception modules and fine-tuning connected layer. In 8th International Conference on Wireless Communications & Signal Processing, WCSP 2016, Yangzhou, China, October 13-15, 2016. IEEE, 1–6. https://doi.org/10.1109/WCSP.2016.7752571
- [11] Erik Learned-Miller, Gary B. Huang, Aruni RoyChowdhury, Haoxiang Li, and Gang Hua. 2016. Labeled Faces in the Wild: A Survey. Advances in Face Detection and Facial Image Analysis (2016), 189–248. https://doi.org/10.1007/ 978-3-319-25958-1_8
- [12] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. 2012. Discovering important people and objects for egocentric video summarization. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012. IEEE Computer Society, 1346–1353. https://doi.org/10.1109/CVPR.2012.6247820
- [13] Jana Machajdik and Allan Hanbury. 2010. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the 18th International Conference on Multimedia 2010, Firenze, Italy, October 25-29, 2010,* Alberto Del Bimbo, Shih-Fu Chang, and Arnold W. M. Smeulders (Eds.). ACM, 83–92. https://doi.org/10.1145/1873951.1873965
- [14] Naila Murray, Luca Marchesotti, and Florent Perronnin. 2012. AVA: A large-scale database for aesthetic visual analysis. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012. IEEE Computer Society, 2408–2415. https://doi.org/10.1109/CVPR.2012.6247954
- [15] Hongwei Ng and Stefan Winkler. 2014. A data-driven approach to cleaning large face datasets. In 2014 IEEE International Conference on Image Processing, ICIP 2014, Paris, France, October 27-30, 2014. IEEE, 343–347. https://doi.org/10.1109/ICIP. 2014.7025068
- [16] Kuang-Chern Ng, Aun Neow Poo, and Marcelo H. Ang. 2001. Practical Issues in Pixel-Based Autofocusing for Machine Vision. In Proceedings of the 2001 IEEE International Conference on Robotics and Automation, ICRA 2001, May 21-26, 2001, Seoul, Korea. IEEE, 2791–2796. https://doi.org/10.1109/ROBOT.2001.933045
- [17] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. 2015. Deep Face Recognition. In Proceedings of the British Machine Vision Conference 2015, BMVC 2015, Swansea, UK, September 7-10, 2015, Xianghua Xie, Mark W. Jones, and Gary K. L. Tam (Eds.). BMVA Press, 41.1–41.12. https://doi.org/10.5244/C.29.41
- [18] Miriam Redi, Nikhil Rasiwasia, Gaurav Aggarwal, and Alejandro Jaimes. 2015. The beauty of capturing faces: Rating the quality of digital portraits. In 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition, FG 2015, Ljubljana, Slovenia, May 4-8, 2015. IEEE Computer Society, 1-8. https://doi.org/10.1109/FG.2015.7163086
- [19] Rossano Schifanella, Miriam Redi, and Luca Maria Aiello. 2015. An Image Is Worth More than a Thousand Favorites: Surfacing the Hidden Beauty of Flickr Pictures. In Proceedings of the Ninth International Conference on Web and Social Media, ICWSM 2015, University of Oxford, Oxford, UK, May 26-29, 2015, Meeyoung Cha, Cecilia Mascolo, and Christian Sandvig (Eds.). AAAI Press, 397–406. http: //www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10547
- [20] Hamid R. Sheikh, Zhou Wang, and Alan C. Bovik. 2002. No-reference perceptual quality assessment of JPEG compressed images. In Proceedings of the 2002 International Conference on Image Processing, ICIP 2002, Rochester, New York, USA, September 22-25, 2002. IEEE, 477–480. https://doi.org/10.1109/ICIP.2002.1038064
- [21] Yale Song, Miriam Redi, Jordi Vallmitjana, and Alejandro Jaimes. 2016. To Click or Not To Click: Automatic Selection of Beautiful Thumbnails from Videos. In Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016, Snehasis Mukhopadhyay, ChengXiang Zhai, Elisa Bertino, Fabio Crestani, Javed Mostafa, Jie Tang, Luo Si, Xiaofang Zhou, Yi Chang, Yunyao Li, and Parikshit Sondhi (Eds.). ACM, 659–668. https://doi.org/10.1145/2983323.2983349
- [22] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. 2015. TVSum: Summarizing web videos using titles. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015.* IEEE Computer Society, 5179–5187. https://doi.org/10.1109/CVPR.2015.7299154
- [23] Pamela H. Vance. 1995. Knapsack Problems: Algorithms and Computer Implementations (S. Martello and P. Toth). SIAM Rev. 35, 4 (1993), 684–685. https://doi.org/10.1137/1035174
- [24] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. 2016. Video Summarization with Long Short-Term Memory. In Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VII (Lecture Notes in Computer Science), Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.), Vol. 9911. Springer, 766–782. https: //doi.org/10.1007/978-3-319-46478-7_47
- [25] Ke Zhou, Miriam Redi, Andrew Haines, and Mounia Lalmas. 2016. Predicting Preclick Quality for Native Advertisements. In Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016, Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao (Eds.). ACM, 299–310. https://doi.org/10.1145/2872427.2883053