Query-adaptive Video Summarization via Quality-aware Relevance Estimation

Arun Balajee Vasudevan^{*†}, Michael Gygli^{*†‡}, Anna Volokitin[†], Luc Van Gool^{†§}

[†]ETH Zurich

[§]KU Leuven

[‡]Gifs.com

{arunv,gygli,anna.volokitin,vangool}@vision.ee.ethz.ch

ABSTRACT

Although the problem of automatic video summarization has recently received a lot of attention, the problem of creating a video summary that also highlights elements relevant to a search query has been less studied. We address this problem by posing query-relevant summarization as a video frame subset selection problem, which lets us optimise for summaries which are simultaneously diverse, representative of the entire video, and relevant to a text query. We quantify relevance by measuring the distance between frames and queries in a common textual-visual semantic embedding space induced by a neural network. In addition, we extend the model to capture query-independent properties, such as frame quality. We compare our method against previous state of the art on textual-visual embeddings for thumbnail selection and show that our model outperforms them on relevance prediction. Furthermore, we introduce a new dataset, annotated with diversity and query-specific relevance labels. On this dataset, we train and test our complete model for video summarization and show that it outperforms standard baselines such as Maximal Marginal Relevance.

1 INTRODUCTION

Video recording devices have become omnipresent. Most of the videos taken with smartphones, surveillance cameras and wearable cameras are recorded with a *capture first, filter later* mentality. However, most raw videos never end up getting curated and remain too long, shaky, redundant and boring to watch. This raises new challenges in searching both within and across videos.

The problem of making videos content more accessible has spurred research in automatic tagging [2, 39, 51] and video summarization [1, 15, 26, 27, 31, 36, 49, 57, 69]. In automatic tagging, the goal is to predict meta-data in form of tags, which makes videos searchable via text queries. Video summarization, on the other hand, aims at making videos more accessible by reducing them to a few interesting and representative frames [26, 31] or shots [15, 56].

This paper combines the goals of summarising videos and makes them searchable with text. Specifically, we propose a novel method that generates video summaries adapted to a text query (See Fig. 1).

^{*}Authors contributed equally

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 the author(s) 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.

MM'17, October 23–27, 2017, Mountain View, CA, USA.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4906-2/17/10...\$15.00

DOI: <https://doi.org/10.1145/3123266.3123297>

GoPro: Heather Jackson - An Ironman

Figure 1: Our query-adaptive video summarization model picks frames that are relevant to the query while also giving a sense of entire video. We want to summarise a video of an *ironman* competition, in which participants swim, bike and run. Query-adapted summaries are representative by showing all three sports, while placing more focus on the frames matching the query.

Our approach improves previous works in the area of textual-visual embeddings [28, 34] and proposes an extension of an existing video summarization method using submodular mixtures [15] for creating summaries that are query-adaptive.

Our method for creating query-relevant summaries consists of two parts. We first develop a relevance model which allows us to rank frames of a video according to their relevance given a text query. Relevance is computed as the sum of the cosine similarity between embeddings of frames and text queries in a learned visual-semantic embedding space and a query-independent term. While the embedding captures semantic similarity between video frames and text queries, the query-independent term predicts relevance based on the quality, composition and the interestingness of the content itself. We train this model on a large dataset of image search data [20] and our newly introduced Relevance and Diversity dataset (Section 5). The second part of the summarization system is a framework for optimising the selected set of frames not only for relevance, but also for representativeness and diversity using a submodular mixture of objectives. Figure 2 shows an overview of our complete pipeline. We make the following contributions:

- Several improvements on learning a textual-visual embedding for thumbnail selection compared to the work by Liu *et al.* [34]. These include better alignment of the learning objective to the task at test time and modeling the text queries using LSTMs, fetching significant performance gains.
- A way to model semantic similarity and quality aspects of frames jointly, leading to better performance compared to using the similarity to text queries only.

- We adapt the submodular mixtures model for video summarization by Gygli *et al.* [15] to create query-adaptive and diverse summaries using our frame-based relevance model.
- A new video thumbnail dataset providing query relevance and diversity labels. As the judgements are subjective, we collect multiple annotations per video and analyse the consistency of the obtained labelling.

2 RELATED WORK

The goal of video summarization is to select a subset of frames that gives a user an idea of the video's content at a glance [60]. To find informative frames for this task, two dominant approaches exist: (i) modelling generic frame interestingness [17, 31] or (ii) using additional information such as the video title or a text query to find relevant frames [33, 34, 56]. In this work we combine the two into one model and make several contributions for query-adaptive relevance prediction. Such models are related to automatic tagging [2, 39, 51], textual-visual embeddings [13, 34, 54] and image description [3, 5, 6, 9, 11, 24, 25, 38]. In the following we discuss approaches for video summarization, generic interestingness prediction models and previous works for obtaining embeddings.

Video summarization. Video summarization methods can be broadly classified into abstractive and extractive approaches. Abstractive or compositional approaches transform the initial video into a more compact and appealing representation, e.g. hyper-lapses [29], montages [58] or video synopses [50]. The goal of extractive methods is instead to select an informative subset of keyframes [26, 27, 31, 63] or video segments [15, 36] from the initial video. Our method is extractive. Extractive methods need to optimise at least two properties of the summary: the quality of the selected frames and their diversity [14, 15, 52]. Sometimes, additional objectives such as temporal uniformity [15] and relevance [52] are also optimised. The simplest approach to obtain a representative and diverse summary is to cluster videos into events and select the best frame per event [7]. More sophisticated approaches jointly optimise for importance and diversity by using determinantal point process (DPPs) [14, 52, 68] or submodular mixtures [15, 32]. Most related to our paper is the work of Sharghi *et al.* [52], who present an approach for query-adaptive video summarization using DPPs. Their method however limits to a small, fixed set of concepts such as *car* or *flower*. The authors leave handling of unconstrained queries, as in our approach, for future work. In this work, we formulate video summarization as a maximisation problem over a set of submodular functions, following [15].

Frame quality/interestingness. Most methods that predict frame interestingness are based on supervised learning. The prediction problem can be formulated as a classification [49], regression [31, 66], or, as is now most common, as a ranking problem [17, 57, 59, 65]. To simplify the task, some approaches assume the domain of the video given and train a model for each domain [49, 57, 65].

An alternative approach based on unsupervised learning, proposed by Xiong *et al.* [64], detects “snap points” by using a web image prior. Their model considers frames suitable as keyframes if the composition of the frames matches the composition of the

web images, regardless of the frame content. Our approach is partially inspired by this work in that it predicts relevance even in the absence of a query, but relies on supervised learning.

Unconstrained Textual-visual models. Several methods exist that can retrieve images given unconstrained text or vice versa [9, 11, 13, 18, 24, 25, 38]. These typically project both modalities into a joint embedding space [13], where semantic similarity can be compared using a measure like cosine similarity. Word2vec [41] and GloVe [48] are popular choices to obtain the embeddings of text. Deep image features are then mapped to the same space via a learned projection. Once both modalities are in the same space, they may be easily compared [13]. A multi-modal semantic embedding space is often used by Zero-shot learning approaches [13, 23, 47] to predict test labels which are unseen in the training. Habibi *et al.* [18], in the same spirit, propose zero-shot recognition of events in videos by learning a video representation that aligns text, audio and video features. Similarly, Liu *et al.* [34] use textual-visual embeddings for video thumbnail selection. Our relevance model is based on Liu *et al.* [34], but we provide several important improvements. (i) Rather than keeping the word representation fixed, we jointly optimise the word and image projection. (ii) Instead of embedding each word separately, we train an LSTM model that combines a complete query into one single embedding vector, thus it even learns multi-word combinations such as *visit to lake* and *Star Wars movie*. (iii) In contrast to Liu *et al.* [34], we directly optimise the target objective. Our experiments show that these changes lead to significantly better performance in predicting relevant thumbnails.

3 METHOD FOR RELEVANCE PREDICTION

The goal of this work is to introduce a method to automatically select a set of video thumbnails that are both relevant with respect to a query, but also diverse enough to represent the video. To later optimise relevance and diversity jointly, we first need a way to evaluate the relevance of frames.

Our relevance model learns a projection of video frames v and text queries t into the same embedding space. We denote the projection of t and v as $\mathbf{t}$ and $\mathbf{v}$, respectively. Once trained, the relevance of a frame v given a query t can be estimated via some similarity measure. As [13], we use the cosine similarity

$$s(\mathbf{t}, \mathbf{v}) = \frac{\mathbf{t} \cdot \mathbf{v}}{\|\mathbf{t}\| \|\mathbf{v}\|}. \quad (1)$$

While this lets us assess the semantic relevance of a frame *w.r.t.* a query, it is also possible to make a prediction on the suitability as thumbnails *a priori*, based on the frame quality, composition, *etc.* [64]. Thus, we propose to extend above notion of relevance and model the quality aspects of thumbnails explicitly by computing the final relevance as the sum of the embedding similarity and the query-independent frame quality term, *i.e.*

$$r(t, v) = s(\mathbf{t}, \mathbf{v}) + q_v, \quad (2)$$

where q_v is a query-independent score determining the suitability of v as a thumbnail, based on the quality of a frame.

In the following, we investigate how to formulate the task of obtaining the embeddings $\mathbf{t}$ and $\mathbf{v}$, as well as q_v .

Figure 2: Overview of our approach. We show how a summary is created from an example video and the query *Cooking channel*. We obtain a query adaptive summary by selecting a set of keyframes from the video using our quality-aware relevance model and submodular mixtures, as explained in Sec. 3 and 4.

3.1 Training objective

Intuitively, our model should be able to answer “What is the best thumbnail for this query?”. Thus, the problem of picking the best thumbnail for a video is naturally formulated as a ranking problem. We desire that the embedding vectors of a query and frame that are a good match are more similar than ones of the same query and a non-relevant frame*. Thus, our model should learn to satisfy the rank constraint that given a query t , the relevance score of the relevant frame v^+ is higher than the relevance score of the irrelevant frame v^- :

$$r(t, v^+) > r(t, v^-). \quad (3)$$

Alternatively, we can train the model by requiring that both the similarity score and the quality score of the relevant frame are higher than for the irrelevant frame explicitly, rather than imposing a constraint only on their sum, as above. In this case we would be imposing the two following constraints:

$$\begin{aligned} s(\mathbf{t}, \mathbf{v}^+) &> s(\mathbf{t}, \mathbf{v}^-) \\ q_{v^+} &> q_{v^-}. \end{aligned} \quad (4)$$

Experimentally, we find that training with these explicit constraints leads to slightly improved performance (See Tab. 1).

In order to impose these constraints and train the model, we define the loss as

$$\begin{aligned} \text{loss}(t, v^+, v^-) &= l_p(\max(0, \gamma - s(\mathbf{t}, \mathbf{v}^+) + s(\mathbf{t}, \mathbf{v}^-))) \\ &+ l_p(\max(0, \gamma - q_{v^+} + q_{v^-})), \end{aligned} \quad (5)$$

where l_p is a cost function and γ is a margin parameter. We follow [17] and use a Huber loss for l_p , i.e. the robust version of an l_2 loss. Next, we describe how to parametrize the $\mathbf{t}$, $\mathbf{v}$ and q_v , so that they can be learned.

3.2 Text and Frame Representation

We use a convolutional neural network (CNN) for predicting $\mathbf{v}$ and q_v , while $\mathbf{t}$ is obtained via a recurrent neural network. To jointly learn the parameters of these networks, we use a Siamese ranking network, trained with triplets of (t, v^+, v^-) where the weights

*Liu *et al.* [34] does the inverse. It poses the problem as learning to assign a higher similarity to corresponding frame and query than to the same frame and a random query. Thus, the model learns to answer the question “what is a good query for this image?”.

for the subnets predicting v^+ and v^- are shared. We provide the model architecture in supplementary material. We now describe the textual representation $\mathbf{t}$ and the image representations $\mathbf{v}$ and q_v in more detail.

Textual representation. As a feature representation $\mathbf{t}$ of the textual query t , we first project each word of the query into a 300-dimensional semantic space using the word2vec model [42], which is trained on GoogleNews dataset. We fine-tune the word2vec model using the unique queries from the Bing Clickture dataset [20] as sentences. Then, we encode the individual word representations into a single fixed-length embedding using an LSTM [19]. We use a many-to-one prediction, where the model outputs a fixed length output at the final time-step. This allows us to emphasize visually informative words and handle phrases.

Image representation. To represent the image, we leverage the feature representations of a pre-trained VGG-19 network [53] on ImageNet [8]. We replace the softmax layer(1000 nodes) of VGG-19 network with a linear layer M with 301 dimensions. The first 300 dimensions are used as the embedding $\mathbf{v}$, while the last dimension represents the quality score q_v .

4 SUMMARIZATION MODEL

We use the framework of submodular optimization to create summaries that take into account multiple objectives [32]. In this framework, summarization is posed as the problem of selecting a subset (in our case, of frames) $\mathbf{y}^*$ that maximizes a linear combination of submodular objective functions $\mathbf{f}(\mathbf{x}_{\mathcal{V}}, \mathbf{y}) = [f_1(\mathbf{x}_{\mathcal{V}}, \mathbf{y}), \dots, f_n(\mathbf{x}_{\mathcal{V}}, \mathbf{y})]^T$. Specifically,

$$\mathbf{y}^* = \arg \max_{\mathbf{y} \in \mathcal{Y}_{\mathcal{V}}} \mathbf{w}^T \mathbf{f}(\mathbf{x}_{\mathcal{V}}, \mathbf{y}), \quad (6)$$

where $\mathcal{Y}_{\mathcal{V}}$ denote the set of all possible solutions $\mathbf{y}$ and $\mathbf{x}_{\mathcal{V}}$ the features of video $\mathcal{V}$. In this work, we assume that the cardinality $|\mathbf{y}|$ is fixed to some value k (we use $k = 5$ in our experiments).

For non-negative weights $\mathbf{w}$, the objective in Eq. (6) is submodular [30], meaning that it can be optimized near-optimally in an efficient way using a greedy algorithm with lazy evaluations [43, 46].

Objective functions. We choose a small set of objective functions, each capturing different aspects of the summary.

- (1) Query similarity $f(\cdot, \cdot) = \sum_{v \in y} s(\mathbf{t}, \mathbf{v})$ where $\mathbf{t}$ is the query embedding, $\mathbf{v}$ is frame embedding and $s(\cdot, \cdot)$ denotes the cosine similarity defined in Eq. (1).
- (2) Quality score $f(\cdot, \cdot) = \sum_{v \in y} q_v$, where q_v represents score that is based on the quality of v as a thumbnail. This model scores the image relevance in a query-independent manner based on properties such as contrast, composition, etc.
- (3) Diversity of the elements in the summary
 $f(\mathbf{x}_y, \mathbf{y}) = \sum_{i \in y} \min_{j < i} D_{x_y}(i, j)$, according to some dissimilarity measure D . We use the Euclidean distance in of the FC2 features of the VGG-19 network for $D^\dagger$.
- (4) Representativeness [15]. This objective favors selecting the medoid frames of a video, such that the visually frequent frames in the video are represented in the summary.

Weight learning. To learn the weights $\mathbf{w}$ in Eq. (6), ground truth summaries for query-video pairs are required. Previous methods typically only optimized for relevance [34] or used small datasets with limited vocabularies [52]. Thus, to be able to train our model, we collected a new dataset with relevance and diversity annotations, which we introduce in the next Section.

If relevance and diversity labels are known, we can estimate the optimal mixing weights of the submodular functions through subgradient descent [32]. In order to directly optimize for the F1-score used at test time, we use a locally modular approximation based on the procedure of [45] and optimize the weights using AdaGrad [10].

5 RELEVANCE AND DIVERSITY DATASET (RAD)

We collected a dataset with query relevance and diversity annotation to let us train and evaluate query-relevant summaries. Our dataset consists of 200 videos, each of which was retrieved given a different query.

Using Amazon Mechanical Turk (AMT) we first annotate the video frames with query relevance labels, and then partition the frames into clusters according to visual similarity. These kind of labels were used previously in the MediaEval diverse social images challenge [21] and enabled evaluation of the automatic methods for creating relevant and diverse summaries.

To select a representative sample of queries and videos for the dataset, we used the following procedure: We take the top YouTube queries between 2008 and 2016 from 22 different categories as seed queries[‡]. These queries are typically rather short and generic concepts, so to obtain longer, more realistic queries we use YouTube auto-complete to suggest phrases. Using this approach we collect 200 queries. Some examples are *brock lesnar vs big show*, *taylor swift out of the woods*, etc. For each query, we take the top video result with a duration of 2 to 3 minutes.

To annotate the videos, we set up two consecutive tasks on AMT. All videos are sampled at one frame per second. In the first task, a worker is asked to label each frame with its relevance *w.r.t.* the given query. Options for answers are “Very Good”, “Good”,

“Not good” and “Trash”, where trash indicates that the frame is both irrelevant and low-quality (e.g. blurred, bad contrast, etc.). After annotating the relevance, the worker is asked to distribute the frames into clusters according to their visual similarity. We obtain one clustering per worker, where each clustering consists of mutually exclusive subsets of video frames as clusters. The number of clusters in the clustering is chosen by the worker. Each video is annotated by 5 different people and a total of 48 subjects participated in the annotation. To ensure high-quality annotations, we defined a qualification task, where we check the results manually to ensure the workers provide good annotations. Only workers who pass this test are allowed to take further assignments.

5.1 Analysis

We now analyse the two kinds of annotations obtained through this procedure and describe how we merge these annotations into one set of ground truth labels per video.

Label distributions. The distribution of relevance labels is “Very Good”: 17.55%, “Good”: 57.40%, “Not good”: 12.31% and “Trash”: 12.72%. The minimum, maximum and mean number of clusters per video are 4.9, 25.2 and 13.4 respectively over all videos of RAD.

Relevance annotation consistency. Given the inherent subjectivity of the task, we want to know whether annotators agree with each other about the query relevance of frames. To do this, we follow previous work [16, 22, 62] and compute the Spearman rank correlation (ρ) between the relevance scores of different subjects, splitting five annotations of each video into two groups of two and three raters each. We take all split combination to find mean ρ for a video.

Our dataset has an average correlation of $\rho = 0.73$ over all videos, where 1 is a perfect correlation while 0 would indicate no consistency in the scores. On the related task of event-specific image importance, using five annotators, consistency is only $\rho = 0.4$ [62]. Thus, we can be confident that our relevance labels are of high quality.

Cluster consistency. To the best of our knowledge, we are the first to annotate multiple clusterings per video and look into the consistency of multiple annotators. MediaEval, for example, used multiple relevance labels but only one clustering [21]. Various ways of measuring the consistency of clusterings exist, e.g. Variation of Information, Normalised Mutual Information or the Rand index (See Wagner and Wagner [61] for an excellent overview). In the following we propose to use Normalised Mutual Information (NMI), an information theoretic measure [12] which is the ratio of the mutual information between two clusterings ($I(C, C')$) and the sum of entropies of the clusterings ($H(C) + H(C')$):

$$NMI(C, C') = \frac{2 \cdot I(C, C')}{H(C) + H(C')}, \quad (7)$$

We chose NMI over the more recently proposed Variation of Information (VI) [40], as NMI has a fixed range ($[0, 1]$) while still being closely related to VI (see supplementary material).

Our dataset has a cluster consistency of 0.54. Since NMI is 0 if two clusterings are independent and 1 iff they are identical, we see that our annotators have a high degree of agreement.

[†]Derivation of submodularity of this objective is provided in the suppl.

[‡]<https://www.google.com/trends/explore>

Method	Settings			Metrics		
	Cost	LSTM	Quality	HIT@1 VG or G	Spear Corr.	mAP
Random	-	-	-	57.17 ± 1.5	-	0.5780
Loss of Liu <i>et al.</i>	l_1	×	×	68.75	0.186	0.6308
Ours: L1	l_1	×	×	68.09	0.209	0.6348
Ours: Huber	l_{huber}	×	×	68.35	0.279	0.6446
Loss of Liu <i>et al.</i> + LSTM	l_1	✓	×	70.62	0.270	0.6507
Ours: Huber + LSTM	l_{huber}	✓	×	72.63	0.367	0.6685
Ours: Frame quality only Q_{expli}	l_{huber}	×	✓	65.95	0.236	0.6315
Ours: Huber + LSTM + Q_{impli}	l_{huber}	✓	✓	70.76	0.371	0.6657
Ours: Huber + LSTM + Q_{expli}	l_{huber}	✓	✓	74.76	0.376	0.6712

Table 1: Comparison of different model configurations trained on a subset of the Clickture dataset and fine-tuned on our Video Thumbnail dataset (RAD). We report the HIT@1 (fraction of times we select a “Very Good” or “Good” thumbnail), the Spearman correlation of our model predictions with the true candidate thumbnail scores, and mean average precision. The Huber+LSTM+ Q_{expli} model performs best.

Ground truth For evaluation on the test videos, we create a single ground truth annotation for each video. We merge the five relevance annotations as well as the clustering of each query-video pair. For the final ground truth of relevance prediction, we require the labels be either positive or negative for each video frame. We map all “Very Good” labels to 1, “Good” labels to 0.5 and “Not Good” and “Trash” labels to 0. We compute the mean of the five relevance annotation labels and label the frame as positive if the mean is ≥ 0.5 and as negative otherwise.

To merge clustering annotations, we calculate NMI between all pairs of clustering and choose the clustering with the highest mean NMI, *i.e.* the most prototypical cluster. An example of relevance and clustering annotation is provided in Fig. 6.

6 CONFIGURATION TESTING

Before comparing our proposed relevance model against state of the art in Sec. 7, we first analyze our model performance using different objectives, cost functions and text representation. For evaluation, we use Query-dependent Thumbnail Selection Dataset (QTS) provided by [34]. The dataset contains 20 candidate thumbnails for each video, each of which is labeled one of the five: Very Good (VG), Good (G), Fair (F), Bad (B), or Very Bad (VB). We evaluate on the available 749 query-video pairs. To transform the categorical labels to numerical values, we use the same mapping as [34].

Evaluation metrics. As evaluation metrics, we are using HIT@1 and mean Average Precision (mAP) as reported and defined in Liu *et al.* [34], as well as the Spearman’s Rank Correlation. HIT@1 is computed as the hit ratio for the highest ranked thumbnail.

Training dataset. For training, we use two datasets: (i) the Bing Clickture dataset [20] and (ii) the RAD dataset (Sec. 5). Clickture is a large dataset consisting of queries and retrieved images from Bing Image search. The annotation is in form of triplets (K, Q, C) meaning that the image K was clicked C times in the search results of the query Q . This dataset is well suited for training our relevance model, since our task is the retrieval of relevant keyframes from a video, given a text query. It is, however, from the image and not the video domain. Thus, we additionally fine-tune the models on the

complete RAD dataset consisting of 200 query-video pairs. From each query-video pair, we sample an equal number of positive and negative frames to give equal weight to each video. In total, we use 0.5M triplets (as in Sec. 3.2) from the Clickture and 14K triplets from the RAD for training.

Implementation details. We preprocess the images as in [53]. We truncate the number of words in the query at 14, as a trade-off between the mean and maximum query length in Clickture dataset (5 and 26 respectively) [44]. We set the margin parameter γ in the loss in Eq. (5) to 1 and the tradeoff parameter δ for the Huber loss to 1.5 as in [17]. The LSTM consists of a hidden layer with 512 units. We train the parameters of the LSTM and projection layer M using stochastic gradient descent with adaptive weight updates (AdaGrad) [10]. We add an l_2 penalty on the weights, with a λ of 10^{-3} . We train for 20 epochs using minibatches of 128 triplets.

6.1 Tested components

We discuss three important components of our model next.

Objective. We compare our proposed training objective to that of Liu *et al.* [34]. Their model is trained to rank a positive query higher than a negative query given a fixed frame. In contrast, our method is trained to rank a positive frame higher than a negative frame given a fixed query.

Cost function. We also investigate the importance of modeling frame quality. In particular, we compare different cost functions. (i) We enforce two ranking constraints: one for the quality term and one for the embedding similarity, as in Eq.(4) (Q_{expli}), (ii) We sum the quality and similarity term into one output score, for which we enforce the rank constraint, as in Eq.(3) (Q_{impli}) or (iii) we don’t model quality at all.

Text representation. As mentioned in Sec. 3.2, we represent the words of the query using word vectors. To combine the individual word representations into single vector, we investigate two approaches: (i) averaging the word embedding vectors and (ii) using an LSTM model that learns to combine the individual word embeddings.

Method	HIT @ 1		Spear. ρ	mAP
	VG	VG or G		
QUERIES				
Liu <i>et al.</i> [34]	40.625	73.83	0.122	0.629
TITLES				
QAR without Q_{expli}	36.71	72.63	0.367	0.6685
QAR (Ours)	38.86	74.76	0.376	0.6712

Table 2: Comparison of thumbnail selection performance against the state of the art, on the QTS evaluation dataset. Note that [34] uses queries for their method which are not publicly available (see text).

Figure 3: Precision-Recall curve of QTS Evaluation dataset for different methods.

6.2 Results

We show the results of our detailed experiments in Tab. 1. They give insights on several important points.

Text representation. Modeling queries with an LSTM, rather than averaging the individual word representations, improves performance significantly. This is not surprising, as this model can learn to ignore words that are not visually informative (*e.g.* 2014).

Objective and Cost function. The analysis shows that training with our objective leads to better performance compared to using the objective of Liu *et al.* [34]. This can be explained with the properties of videos, which typically contain many frames that are low-quality or not visually informative [55]. Thus, formulating the thumbnail task in a way that the model can learn about these quality aspects is beneficial. Using the appropriate triplets for training boosts performance substantially (correlation with the loss of Liu *et al.* [34] + LSTM: 0.270, Ours: Huber + LSTM 0.367). When including a quality term in the model, performance improves further, where an explicit loss performs slightly better (Ours: Huber + LSTM + Q_{expli} in Tab. 1).

Somewhat surprisingly, modeling quality alone already outperforms Liu *et al.* [34] in terms of mAP, despite not using any textual information. Quality adds a significant boost to performance in the

Method	HIT@1	Spear. ρ	mAP
No TEXTUAL INPUT			
Random	66.6 $\pm$ 3.5	0.0	0.674
Video2GIF [17]	67.0	0.167	0.708
Ours: Frame quality Q_{expli}	69.0	0.135	0.749
TITLES			
Liu <i>et al.</i> [34] + LSTM	70.0	0.134	0.731
QAR without Q_{expli}	70.0	0.182	0.743
QAR (Ours)	71.0	0.221	0.760
QUERIES			
Liu <i>et al.</i> [34] + LSTM	72.0	0.204	0.730
QAR without Q_{expli}	76.0	0.268	0.752
QAR (Ours)	72.0	0.264	0.769

Table 3: Performance of our relevance models on the RAD dataset in comparison with previous methods.

Figure 4: Recall-Precision curve of the RAD testset for different methods. Our method (Blue) performs high in terms of mAP.

video domain. Interestingly, this is different in the image domain, due to the difference in quality statistics. Images returned by a search engine are mostly of good quality, thus explicitly accounting for it does not improve performance (see supplementary material).

To conclude, we see that the better alignment of the objective to the keyframe retrieval task, the addition of an LSTM and modeling quality of the thumbnails improves performance. Together, they provide a substantial improvement compared to Liu *et al.*'s model. Our method achieves an absolute improvement of 6.01% in HIT@1, 4.04% in mAP, and an improvement in correlation from 0.186 to 0.376. These gains are even more significant when we consider the possible ranges of these metrics. *e.g.* for Spearman correlation, human agreement is at 0.73 on the RAD dataset (*c.f.* Sec. 5.1), thus providing an upper bound. Similarly, HIT@1 and mAP have small effective ranges given their high scores for a random model.

Figure 5: Qualitative Results of top ranked keyframes on RAD. a) Liu *et al.* b) Video2GIF c) Our model (from left). Video titles are shown on the left. Ground truth relevance labels are shown in Blue. P=Positive, N=Negative.

7 EXPERIMENTS

In the previous section, we have determined that our objective, embedding queries with an LSTM and explicitly modelling quality performs best. We call this model QAR (Quality-Aware Relevance) in the following and compare against state-of-the-art(s-o-a) models on the QTS and RAD datasets. We also evaluate the full summarization model on RAD. For these experiments, we split RAD into 100 videos for training, 50 for validation and 50 for testing.

Evaluation metrics. For relevance we use the same metrics as in Sec. 6. To evaluate video summaries on RAD, we additionally use F1 scores. The F1 score is the harmonic mean of precision of relevance prediction and cluster recall [21]. It is high, if a method selects relevant frames from diverse clusters.

7.1 Evaluating the Relevance Model

We evaluate our model (QAR) and compare it to Liu *et al.* [34] and Video2GIF [17].

Query-dependent Thumbnail Selection Dataset (QTS) [34]

We compare against the s-o-a on the QTS evaluation dataset in Tab. 2. We report the performance of Liu *et al.* [34] from their paper. Note, however, that the results are not directly comparable, as they use query-video pairs for predicting relevance, while only the titles are shared publicly. Thus, we use the titles instead, which is an important difference. Relevance is annotated with respect to the queries, which often differ from the video titles. We compare the re-implementation of [34] using titles in detail in Tab. 1.

Encouragingly, our model performs well even when just using the titles and outperforms them on most metrics. It improves mAP by 4.22% over [34] and correlation by a margin of 0.254 (c.f. Table 2). Figure 3 shows the precision-recall curve for the experiment. As can be seen QAR outperforms [34] for all recall ratios. To better understand the effects of using titles or queries, we quantify the value of the two on the RAD dataset below.

Method				< PR >	< CR >	< F1 >
Similarity	Diversity	Quality	Repr			
–	–	–	✓	0.654	0.817	0.672
–	–	✓	–	0.671	0.542	0.522
–	✓	–	–	0.575	0.808	0.629
✓	–	–	–	0.763	0.550	0.578
✓	–	✓	–	0.775	0.563	0.594
MMR [4]						
✓ (33%)	✓ (66%)	–	–	0.692	0.825	0.716
Hecate [55]				0.708	0.787	0.713
Ours						
✓ (45%)	✓ (43%)	✓ (2%)	✓ (10%)	0.704	0.825	0.721
Upper bound				0.938	0.925	0.928

Table 4: Performance of summarization methods on the RAD dataset. *Repr* means Representativeness. ✓ and – depict whether an objective was used or not. MMR and ours learn their corresponding weights. Percentage in parentheses the normalized learnt weights. Upper bound refers to the best possible performance, obtained using the ground truth annotations of RAD.

Our dataset (RAD) We also evaluate our model on the RAD test set (Tab. 3). QAR (ours) significantly outperforms the previous s-o-a of [17, 34], even when augmenting Liu *et al.* [34] with an LSTM. QAR improves mAP by 2.9% when using TITLES and 3.9% when using QUERIES over our implementation of Liu *et al.* [34]+LSTM.

We also see that modeling quality leads to significant gains in terms of mAP when using TITLES or QUERIES (+1.7% in both cases). HIT@1 for query relevance, however, is lower when including quality. We believe that the reason for this is that when the query is given, the textual-visual similarity is a more reliable signal to determine the single best keyframe. While including quality improves the overall ranking on mAP, it is solely based on appearance and thus seems to inhibit the fine-grained ranking results at low recall (Fig. 4). However, when only the title is used, the frame quality becomes a stronger predictor for thumbnail selection and improves performance on all metrics. We present some qualitative results of different methods for relevance prediction in Fig. 5.

7.2 Evaluating the Summarization Model

As mentioned in Sec. 4, we use four objectives for our summarization model. Referring to Tab. 4, we use QAR model to get Similarity and Quality scores while Diversity and Representativeness scores are obtained as described in Sec. 4. We compare the performance of our full model with each individual objective, a baseline based on Maximal Marginal Relevance (MMR) [4] and Hecate [55]. MMR greedily builds a set that maximises the weighted sum of two terms: (i) The similarity of the selected elements to a query and (ii) The dissimilarity to previously selected elements. To estimate the similarity to the query we use our own model (QAR without Q_{expli}) and for dissimilarity the diversity as defined in Sec. 4. Finally, we compare it to Hecate, recently introduced in [55]. Hecate estimates frame quality using the stillness of the frame and selects representative and diverse thumbnails by clustering the video with k-means and selecting the highest quality frame from the k largest clusters.

Figure 6: We show video summaries created by Hecate [55], MMR [4], our similarity model and our full summarization approach. The **Green** number on the images depicts the frame number. We plot the ground truth relevance scores, marking the selected frames for the shown methods, and cluster annotations over the video in the bottom two rows. For cluster annotation, each color represents a unique cluster. Additional examples are provided in supplementary.

Results Quantitative results are shown in Tab. 4, while Fig. 6 shows qualitative results. As can be seen, combining all objectives with our model works best. It outperforms all single objectives, as well as the MMR [4] baseline, even though MMR also uses our well-performing similarity estimation. Similarity alone has the highest precision, but tends to pick frames that are visually similar (*c.f.* Fig. 6), thus resulting in low cluster recall. Diversification objectives (diversity and representativeness) have a high cluster recall, but the frames are less relevant. Somewhat surprisingly, Hecate [55] is a relatively strong baseline. In particular, it performs well in terms of relevance, despite using a simple quality score. This further highlights the importance of quality for the thumbnail selection task. It also indicates that the used VGG-19 architecture might be suboptimal for predicting quality. CNNs for classification use small input resolutions, thus making it difficult to predict quality aspects such as blur. Finding better architectures for that task is actively researched, *e.g.* [35, 37], and might be used to improve our method.

When analysing the learned weights (*c.f.* Tab. 4) we find that the similarity prediction is the most important objective, which matches our expectations. Quality gets a lower, but non-zero weight, thus showing that it provides information that is complementary to query-similarity. Thus, it helps predicting the relevance of a frame. The reader should however be aware that differences in the variance of the objectives can affect the weights learned. Thus, they should be taken with a grain of salt and only be considered tendencies.

8 CONCLUSION

We introduced a new method for query-adaptive video summarization. At its core lies a textual-visual embedding, which lets us select frames relevant to a query. In contrast to earlier works, such as [52, 67], this model allows us to handle unconstrained queries and even full sentences. We proposed and empirically evaluated different improvements over [34], for learning a relevance model. Our empirical evaluation showed that a better training objective, a more sophisticated text model, and explicitly modelling quality leads to significant performance gains. In particular, we showed that quality plays an important role in the absence of high-quality relevance information, such as queries, *i.e.* when only the title can be used. Finally, we introduced a new dataset for thumbnail selection which comes with query-relevance labels and a grouping of the frames according to visual and semantic similarity. On this data, we tested our full summarization framework and showed that it compares favourably to strong baselines such as MMR [4] and [55]. We hope that our new dataset will spur further research in query adaptive video summarization.

9 ACKNOWLEDGEMENTS

This work has been supported by Toyota via the project TRACE-Zurich. We also acknowledge the support by the CHIST-ERA project MUSTER. MG was supported by the European Research Council under the project VarCity (#273940).

REFERENCES

- [1] I Arev, HS Park, and Yaser Sheikh. 2014. Automatic editing of footage from multiple social cameras. *ACM Transactions on Graphics (TOG)* (2014).
- [2] Lamberto Ballan, Marco Bertini, Giuseppe Serra, and Alberto Del Bimbo. 2015. A data-driven approach for tag refinement and localization in web videos. *Computer Vision and Image Understanding* (2015).
- [3] Andrei Barbu, Alexander Bridge, Zachary Burchill, Dan Coroian, Sven Dickinson, Sanja Fidler, Aaron Michaux, Sam Mussman, Siddharth Narayanaswamy, Dhaval Salvi, Lara Schmidt, Jiangnan Shangquan, Jeffrey Mark Siskind, Jarrell Waggoner, Song Wang, Jinlian Wei, Yifan Yin, and Zhiqi Zhang. 2012. Video In Sentences Out. *UAI* (2012). arXiv:arXiv:1204.2742v1
- [4] Jaime Carbonell and Jade Goldstein. 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *ACM SIGIR*.
- [5] Xinlei Chen and C Lawrence Zitnick. 2014. Learning a Recurrent Visual Representation for Image Caption Generation. *Proceedings of CoRR* (2014).
- [6] Pradipto Das, Chenliang Xu, Richard F. Doell, and Jason J. Corso. 2013. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. *CVPR* (2013).
- [7] Sandra E. F. de Avila, Ana P. B. Lopes, A. da Luz, and A. de Albuquerque Araújo. 2011. VSUMM: A mechanism designed to produce static video summaries and a novel evaluation method. *Pattern Recognition Letters* 32, 1 (2011).
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- [9] J. Donahue, Lisa A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, T. Darrell, and K. Saenko. 2015. Long-term recurrent convolutional networks for visual recognition and description. *CVPR* (2015).
- [10] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research* 12 (2011).
- [11] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *CVPR*.
- [12] Ana L. N. Fred and Anil K. Jain. 2003. Robust Data Clustering. *CVPR* (2003).
- [13] Andrea Frome, Gs Corrado, and Jonathon Shlens. 2013. Devise: A deep visual-semantic embedding model. *NIPS* (2013). arXiv:arXiv:1312.5650v3
- [14] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. 2014. Diverse sequential subset selection for supervised video summarization. In *NIPS*.
- [15] Michael Gygli, Helmut Grabner, and Luc Van Gool. 2015. Video Summarization by Learning Submodular Mixtures of Objectives. In *CVPR*.
- [16] M Gygli, H Grabner, H Riemenschneider, F Nater, and L Van Gool. 2013. The Interestingness of Images. In *ICCV*.
- [17] Michael Gygli, Yale Song, and Liangliang Cao. 2016. Video2GIF: Automatic Generation of Animated GIFs from Video. *CVPR* (2016).
- [18] Amirhossein Habibian, Thomas Mensink, and Cees GM Snoek. 2016. VideoStory embeddings recognize events when examples are scarce. *TPAMI* (2016).
- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* (1997).
- [20] XS Hua, L Yang, M Ye, K Wang, Y Rui, and J Li. 2013. *Clickture: A large-scale real-world image dataset*. Technical Report. MSR-TR-2013-75.
- [21] Bogdan Ionescu, Adrian Popescu, Mihai Lupu, Alexandru Lucian Ginsca, and Henning Müller. 2015. Retrieving Diverse Social Images at MediaEval 2015: Challenge, Dataset and Evaluation. *MediaEval* (2015).
- [22] Phillip Isola, Devi Parikh, Antonio Torralba, and Aude Oliva. 2011. Understanding the intrinsic memorability of images. In *NIPS*.
- [23] M. Jain, Jan C van Gemert, T. Mensink, and C. GM Snoek. 2015. Objects2action: Classifying and localizing actions without any video example. In *ICCV*.
- [24] Andrej Karpathy, Armand Joulin, and Fei Fei Li. 2014. Deep fragment embeddings for bidirectional image sentence mapping. *NIPS* (2014).
- [25] Andrej Karpathy and Fei Fei Li. 2015. Deep visual-semantic alignments for generating image descriptions. *CVPR* (2015).
- [26] Aditya Khosla, Raffay Hamid, CJ Lin, and Neel Sundaresan. 2013. Large-Scale Video Summarization Using Web-Image Priors. *CVPR* (2013).
- [27] Gunhee Kim, Leonid Sigal, and Eric P Xing. 2014. Joint summarization of large-scale collections of web images and videos for storyline reconstruction. In *CVPR*.
- [28] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539* (2014).
- [29] Johannes Kopf, Michael F Cohen, and Richard Szeliski. 2014. First-person hyper-lapse videos. *ACM Transactions on Graphics* (2014).
- [30] A. Krause and D. Golovin. 2012. Submodular function maximization. (2012).
- [31] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. 2012. Discovering important people and objects for egocentric video summarization. *CVPR* (2012).
- [32] Hui Lin and JA Bilmes. 2012. Learning mixtures of submodular shells with application to document summarization. *arXiv preprint arXiv:1210.4871* (2012).
- [33] Feng Liu, Yuzhen Niu, and Michael Gleicher. 2009. Using Web Photos for Measuring Video Frame Interestingness. *IJCAI* (2009).
- [34] Wu Liu, Tao Mei, Yongdong Zhang, C Che, and Jiebo Luo. 2015. Multi-task deep visual-semantic embedding for video thumbnail selection. *CVPR* (2015).
- [35] Xin Lu, Zhe Lin, Xiaohui Shen, Radomir Mech, and James Z Wang. 2015. Deep multi-patch aggregation network for image style, aesthetics, and quality estimation. In *CVPR*.
- [36] Zheng Lu and Kristen Grauman. 2013. Story-driven summarization for egocentric video. *CVPR* (2013).
- [37] Long Mai, Hailin Jin, and Feng Liu. 2016. Composition-preserving deep photo aesthetics assessment. In *CVPR*.
- [38] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L Yuille. 2014. Explain images with multimodal recurrent neural networks. *arXiv preprint arXiv:1410.1090* (2014).
- [39] Masoud Mazloom, Xirong Li, and Cees Snoek. 2016. TagBook: A Semantic Video Representation without Supervision for Event Detection. *IEEE Transactions on Multimedia* (2016).
- [40] Marina Meilă. 2003. Comparing clusterings by the variation of information. In *Learning theory and kernel machines*. Springer.
- [41] T. Mikolov, K. Chen, G. Corrado, and J. Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [42] T Mikolov and J Dean. 2013. Distributed representations of words and phrases and their compositionality. *NIPS* (2013).
- [43] M Minoux. 1978. Accelerated greedy algorithms for maximizing submodular set functions. *Optimization Techniques* (1978).
- [44] Jonas Mueller and Aditya Thyagarajan. 2016. Siamese Recurrent Architectures for Learning Sentence Similarity. In *AAAI*.
- [45] Mukund Narasimhan and Jeff A Bilmes. 2012. A submodular-supermodular procedure with applications to discriminative structure learning. *arXiv preprint arXiv:1207.1404* (2012).
- [46] GL Nemhauser, LA Wolsey, and ML Fisher. 1978. An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming* (1978).
- [47] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, and J. Dean. 2013. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650* (2013).
- [48] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global Vectors for Word Representation. In *EMNLP*.
- [49] Danila Potapov, Matthijs Douze, Zaid Harchaoui, and Cordelia Schmid. 2014. Category-specific video summarization. In *ECCV*. Springer.
- [50] Yael Pritch, Alex Rav-Acha, and Shmuel Peleg. 2008. Nonchronological video synopsis and indexing. *TPAMI* 30, 11 (2008).
- [51] GJ Qi, XS Hua, Y Rui, J Tang, T Mei, and HJ Zhang. 2007. Correlative Multi-Label Video Annotation. (2007).
- [52] Aidean Sharghi, Boqing Gong, and Mubarak Shah. 2016. Query-Focused Extractive Video Summarization. *CoRR abs/1607.05177* (2016).
- [53] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [54] Richard Socher, Andrej Karpathy, Quoc V Le, Christopher D Manning, and Andrew Y Ng. 2014. Grounded compositional semantics for finding and describing images with sentences. *ACL* (2014).
- [55] Y. Song, M. Redi, J. Vallmitjana, and A. Jaimes. 2016. To Click or Not To Click: Automatic Selection of Beautiful Thumbnails from Videos. In *CIKM*. ACM.
- [56] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. 2015. TVSUM: Summarizing web videos using titles. In *CVPR*.
- [57] Min Sun, Ali Farhadi, and Steve Seitz. 2014. Ranking Domain-Specific Highlights by Analyzing Edited Videos. *ECCV* (2014).
- [58] Min Sun, Ali Farhadi, Ben Taskar, and Steve Seitz. 2014. Salient Montages from Unconstrained Videos. *ECCV* (2014).
- [59] Min Sun, Kuo-Hao Zeng, Yenchun Lin, and Farhadi Ali. 2017. Semantic Highlight Retrieval and Term Prediction. *IEEE Transactions on Image Processing* (2017).
- [60] Ba Tu Truong and Svetha Venkatesh. 2007. Video abstraction. *ACM Transactions on Multimedia Computing, Communications, and Applications* (2007).
- [61] Silke Wagner and Dorothea Wagner. 2007. Comparing Clusterings - An Overview. *Graph-Theoretic Concepts in Computer Science* (2007).
- [62] Yufei Wang, Zhe Lin, Xiaohui Shen, Radomir Mech, Gavin Miller, and Garrison W Cottrell. 2016. Event-Specific Image Importance. In *CVPR*.
- [63] Wayne Wolf. 1996. Key frame selection by motion analysis. *Acoustics, Speech, and Signal Processing, ICASSP-96* (1996).
- [64] Bo Xiong and Kristen Grauman. 2014. Detecting snap points in egocentric video with a web photo prior. In *ECCV*.
- [65] Ting Yao, Tao Mei, and Yong Rui. 2016. Highlight detection with pairwise deep ranking for first-person video summarization. In *CVPR*.
- [66] Gloria Zen, Paloma de Juan, Yale Song, and Alejandro Jaimes. 2016. Mouse activity as an indicator of interestingness in video. In *ICMR*.
- [67] Kuo-Hao Zeng, Yen-Chen Lin, Ali Farhadi, and Min Sun. 2016. Semantic highlight retrieval. In *ICIP*.
- [68] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. 2016. Video Summarization with Long Short-term Memory. *ECCV* (2016).
- [69] Bin Zhao and Eric P Xing. 2014. Quasi real-time summarization for consumer videos. In *CVPR*.