

Query-Focused Extractive Video Summarization

Aidean Sharghi^(✉), Boqing Gong, and Mubarak Shah

Center for Research in Computer Vision, University of Central Florida, Orlando, USA
aidean.sharghi@knights.ucf.edu, {bgong,shah}@crcv.ucf.edu

Abstract. Video data is explosively growing. As a result of the “big video data”, intelligent algorithms for automatic video summarization have (re-)emerged as a pressing need. We develop a probabilistic model, Sequential and Hierarchical Determinantal Point Process (SH-DPP), for **query-focused** extractive video summarization. Given a user query and a long video sequence, our algorithm returns a summary by selecting key shots from the video. The decision to include a shot in the summary depends on the shot’s relevance to the user query and importance in the context of the video, jointly. We verify our approach on two densely annotated video datasets. The query-focused video summarization is particularly useful for search engines, e.g., to display snippets of videos.

1 Introduction

Video data is explosively growing as a result of the ubiquitous acquisition capabilities. The videos captured by UAVs and/or drones, from ground surveillance, and by body-worn cameras can easily reach the scale of gigabytes per day. About 300 h of videos are uploaded per minute to Youtube. While the “big video data” is a great source for information discovery, the computational challenges are unparalleled. In such context, intelligent algorithms for automatic video summarization (and retrieval, recognition, etc.) (re-)emerge as a pressing need.

In this paper we focus on *extractive* video summarization, which generates a concise summary of a video by selecting from it key frames or shots¹. The key frames/shots are expected to be (1) individually important—otherwise they should not be selected, and (2) collectively diverse—otherwise one can remove some of them without losing much information. These two principles are employed in most of the existing works on extractive video summarization [4], and yet implemented by different decision choices. Some earlier works define the importance of key frames by low-level appearance and/or motion cues [5–10]. Contextual information of a key frame is often modeled by graphs [11–13]. We

¹ It is also appealing to have the summary as a spatial-temporal synopsis or mosaic *composed* of multiple frames. However, the *compositional* summarization is challenging and has achieved some success in only well-controlled environments [1–3].

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46484-8_1](https://doi.org/10.1007/978-3-319-46484-8_1)) contains supplementary material, which is available to authorized users.

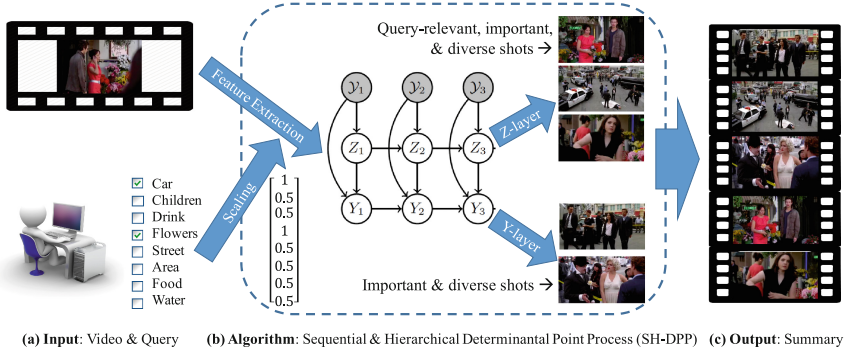


Fig. 1. Query-focused video summarization and our approach to this problem.

note that the system developers play a vital role in this cohort of works; most decisions on how to measure the importance and diversity are handcrafted by the system developers using the low-level cues.

Recently, we see a paradigm shift in some sort: more high-level supervised information is introduced to video summarization than ever before. Rich Web images and videos provide (weak) priors for defining user-oriented importance of the visual content in a video [14–17]. For instance, the CAR images on the Web reveal the canonical views of the cars, which should thus be given special attentions in video summarization. The texts associated with videos are undoubtedly good sources for inferring the semantic importance of video frames [18, 19]. Category-specific and domain-specific video summarization approaches are developed in [20, 21]. Some other high-level factors include gaze [22], interestingness [23], influence [24], tracking of salient objects [25, 26], and so forth.

What are the advantages of leveraging high-level supervised information in video summarization over merely low-level cues? We believe the main advantage is that the system developers are able to better infer the system *users*’ needs. After all, video summarization is a subjective process. Comparing to designing the system from the experts’ own intuitions, it is more desirable to design a system based on the crowd or average users such that the system’s states approach the users’ internal ones, which are often semantic and high-level.

What is the best supervision for a video summarization system? We have seen many types of supervision used in the above-mentioned works, such as Web images, texts, and categories. However, we argue that the best supervision, for the purpose of developing video summarization approaches, is the video summaries directly provided by users. In [27], which is the first supervised video summarization work as far as we know, Gong et al. showed that there exists a high inter-annotator agreement in the summaries of the same videos given by distinct users. They proposed a supervised video summarization model, sequential determinantal point process (seqDPP), and train seqDPP by the “oracle”

summaries that agree the most with different user summaries. Gygli et al. gave another supervised method using submodular functions [28].

From the low-level visual and motion cues to the high-level (indirect) supervised information, and to the (direct) supervised user summaries, video summarization works become more and more **user-oriented**. Though the two principles, importance and diversity, remain the same, the detailed implementation choices have significantly shifted from the system developers’ to the users’; users can essentially teach the system how to summarize videos in [27, 28].

In respect to the recent progress, the goal of this paper is to further advance the user-oriented video summarization by modeling user input, or more precisely user intentions, in the summarization process. Figure 1 illustrates our main idea. We name it **query-focused** (extractive) video summarization, in accordance with the query-focused document summarization [29] in NLP. A query refers to one or more concepts (e.g., CAR, FLOWERS) that are both user-nameable and machine-detectable. More generic queries are left for the future work.

Towards the goal of query-focused video summarization, we develop a probabilistic model, Sequential and Hierarchical Determinantal Point Process (SH-DPP). It has two layers of random variables, each of which serves for subset selection from a ground set of video shots (see Fig. 2). The first layer is mainly used to select the shots relevant to the user queries, and the second layer models the importance of the shots in the context of the videos. We condition the second layer on the first layer so that we can automatically balance the two strengths by learning from user labeled summaries. The determinantal point process (DPP) [30] is employed to account for the diversity of the summary.

A key feature in our work is that the decision to include a video shot in the summary is jointly dependent on the shot’s relevance to the query and representativeness in the video. Instead of handcrafting any criteria, we use SH-DPP to automatically learn from the user summaries (and the corresponding user queries and video sequences). In a sharp contrast to [27, 28] which model average users, our work closely tracks individual users’ intentions from their input queries, and thus has greater potential to satisfy various user needs: distinct personal preferences (e.g., a patient user prefers more detailed and lengthy summaries than an impatient user), different interests over time even about the same video (e.g., a party versus a particular person in the party), etc. Finally, we note that our work is especially useful for search engines to produce snippets of videos.

Our main contribution is on the query-focused video summarization. Querying videos is not only an appealing functionality to the users but also an effective communication channel for the system to capture a user’ intention. Besides, we develop a novel probabilistic model, SH-DPP. Similarly to the sequential DPP (seqDPP) [27], SH-DPP is efficient in modeling extremely lengthy videos and capable of producing summaries on the fly. Additionally, SH-DPP explicitly accounts for the user input queries. Extensive experiments on the UT Egocentric [31] and TV episodes [32] datasets verify the effectiveness of SH-DPP. To our knowledge, our work is the first on query-focused video summarization.

2 Related Work and Background

In this section, we mainly discuss some related works on query-focused document summarization and some earlier works on interactive video summarization in the multimedia community. We will then describe some variations of DPP and contrast them to our SH-DPP.

Query-focused document summarization has been a long-standing track in the Text Retrieval Conference (<http://trec.nist.gov/>) and the Document Understanding Conference (DUC) (<http://duc.nist.gov/>). In DUC 2005, participants were asked to summarize a cluster of documents given a user’s query describing the information needs. Some representative approaches to this problem include BAYESUM [33], FASTSUM [34], and log-likelihood based method [35] among others. Behind the vast research in this topic is the strong motivations by popular search engines and human-machine interactions. However, the counterpart in vision, query-focused video summarization, has not been well formulated yet. We make some preliminary efforts toward it through this work.

Interactive video summarization shares some spirits with our query-focused video summarization. The system in [36] allows users to interactively select some video shots to the summary while the system summarizes the remaining video. In contrast, in our system the users can use concept-based queries to influence the summaries without actually watching the videos. Besides, our approach is supervised and trained by user annotations, not handcrafted by the system developers. There are some other works involving users for thumbnail selection [19] and storyline-based video representation [37]. Our work instead involves user input in the video summarization.

Determinantal point process (DPP) [30] is employed in our SH-DPP to model the diversity in the desired video summaries. We give it a brief overview and also contrast SH-DPP to various DPP models.

Denote by $\mathcal{Y} = \{1, 2, \dots, N\}$ the ground set. A (L-ensemble) DPP defines a discrete probability distribution over a subset selection variable Y ,

$$P(Y = y) = \det(\mathbf{L}_y) / \det(\mathbf{L} + \mathbf{I}), \quad \forall y \subseteq \mathcal{Y}, \quad (1)$$

where \mathbf{I} is an identity matrix, $\mathbf{L} \in \mathbb{S}^{N \times N}$ is a positive semidefinite kernel matrix and is the distribution parameter, and \mathbf{L}_y is a squared sub-matrix with rows and columns corresponding to the indices $y \subseteq \mathcal{Y}$. By default $\det(\mathbf{L}_\emptyset) = 1$.

DPP is good for modeling summarization because it integrates the two principles of individual importance and collective diversity. By the definition (Eq. (1)), the importance of an item is represented by $P(i \in Y) = \mathbf{K}_{ii}$ and the repulsion of any two items is captured by $P(i, j \in Y) = P(i \in Y)P(j \in Y) - \mathbf{K}_{ij}^2$, where $\mathbf{K} = \mathbf{L}(\mathbf{L} + \mathbf{I})^{-1}$. In other words, the model parameter \mathbf{L} is sufficient to describe both the importance and diversity of the items being selected by Y . The readers are referred to Theorem 2.2 in [30] for more derivations.

A vanilla DPP gave rise to state-of-the-art performance on document summarization [38, 39]. Its variation, Markov DPP [40], was used to maintain

the diversity between multiple draws from the ground set. A sequential DPP (seqDPP) [27] was proposed for video summarization. Our SH-DPP brings a hierarchy to seqDPP and uses the first layer to take account of the user queries in the summarization (subset selection) process.

3 Approach

Our approach takes as input a user query q (i.e., concepts) and a long video \mathcal{Y} , and outputs a query-focused short summary $y(q, \mathcal{Y})$,

$$y(q, \mathcal{Y}) \leftarrow \underset{y \subseteq \mathcal{Y}}{\operatorname{argmax}} P(Y = y|q, \mathcal{Y}), \quad (2)$$

which consists of some shots of the video. We desire **four major properties** from the distribution $P(Y = y|q, \mathcal{Y})$. (i) It models the subset selection variable Y . (ii) It promotes diversity among the items selected by Y . (iii) It works efficiently given very long (e.g., egocentric) or endlessly streaming (e.g., surveillance) videos. (iv) It has some mechanism for accepting the user input q . Together, the properties motivate a Sequential and Hierarchical DPP (SH-DPP) as our implementation to $P(Y = y|q, \mathcal{Y})$. As below, we firstly discuss some related methods—especially seqDPP, how they meet some of the **properties** but not all, and then present the details of SH-DPP.

3.1 Sequential DPP (seqDPP) with User Queries

In order to satisfy **properties** (i) and (ii), one can use a vanilla DPP (cf. Eq. (1)) to extract a diverse subset of shots as a video summary. Though this works well for multi-document summarization [38], it is unappealing in our context mainly due to two reasons. First, DPP sees the ground set (i.e., all shots in a video) as a bag, in which the permutation of the items has no effect on the output. In other words, the temporal flow of the video is totally ignored by DPP; it returns the same summary even if the shots are randomly shuffled. Second, the inference (Eq. (2)) cost is extremely high when the video is long, no matter by exhaustive search among all possible subsets $y \subseteq \mathcal{Y}$ or greedy search [30]. We note that the submodular functions also suffer from the same drawbacks [22, 28].

The seqDPP method [27] meets **properties** (i)–(iii) and solves the problems described above. It partitions a video into T consecutive disjoint segments, $\cup_{t=1}^T \mathcal{Y}_t = \mathcal{Y}$, where \mathcal{Y}_t represents a set consisting of only a few shots and stands as the ground set of time step t . The model is defined as follows (see the left panel of Fig. 2 for the graphical model),

$$P_{\text{SEQ}}(Y|\mathcal{Y}) = P(Y_1|\mathcal{Y}_1) \prod_{t=2}^T P(Y_t|Y_{t-1}, \mathcal{Y}_t), \quad \mathcal{Y} = \cup_{t=1}^T \mathcal{Y}_t \quad (3)$$

where $P(Y_t|Y_{t-1}, \mathcal{Y}_t) \propto \det \mathbf{\Omega}_{Y_{t-1} \cup Y_t}$ is a conditional DPP to ensure diversity between the items selected at time step t (by Y_t) and those of the previous time




Fig. 2. The graphical models of seqDPP [27] (left) and our SH-DPP (right).

step (by Y_{t-1}). Similarly to the vanilla DPP (cf. Eq. (1)), here the conditional DPP is also associated with a kernel matrix Ω . In [27], this matrix is parameterized by $\Omega_{ij} = \mathbf{f}_i^T W^T W \mathbf{f}_j$, where \mathbf{f}_i is a feature vector of the i -th video shot and W is learned from the user summaries. Note that the seqDPP summarizer $P_{\text{SEQ}}(Y|\mathcal{Y})$ does not account for any user input. It is learned from “oracle” summaries in the hope of reaching a good compromise between distinct users.

In this paper, we instead aim to infer individual users’ preferences over the video summaries, through the information conveyed by the user queries. To this end, a simple extension to seqDPP is to engineer query-dependent feature vectors $\mathbf{f}(q)$ of the video shots—see Sect. 4.4. We consider this seqDPP variation as our baseline. It is indeed responsive to the queries through the query-dependent features, but it is limited in modeling the query-relevant summaries, in which the importance of a video shot is jointly determined by its relevance to the query and its representativeness in the context. The seqDPP offers no explicit treatment to the two types of interplayed strengths; the user may expect different levels of diversity from the query relevant shots and irrelevant ones, but the single DPP kernel in seqDPP fails to offer such flexibility.

Our SH-DPP possesses all of the four **properties**. It is developed upon seqDPP in order to take advantage of seqDPP’s nice properties (i)–(iii), and yet rectifies its downside (mainly on **property** (iv)) by a two-layer hierarchy.

3.2 Sequential and Hierarchical DPP (SH-DPP)

The right panel of Fig. 2 depicts the graphical model of SH-DPP, reading as,

$$\begin{aligned}
 & P_{\text{SH}}(\{Y_1, Z_1\}, \dots, \{Y_T, Z_T\} | q, \mathcal{Y}) \\
 &= P(Z_1 | q, \mathcal{Y}_1) P(Y_1 | Z_1, \mathcal{Y}_1) \prod_{t=2}^T P(Z_t | q, Z_{t-1}, \mathcal{Y}_t) P(Y_t | Z_t, Y_{t-1}, \mathcal{Y}_t). \quad (4)
 \end{aligned}$$

Query q is omitted from Fig. 2 for clarity. The shaded nodes represent video segments $\{\mathcal{Y}_t\}$ (i.e., consecutive and disjoint shots). We first use the subset selection variables Z_t to select the query-relevant video shots. Note that Z_t will return empty if the segment \mathcal{Y}_t does not contain any visual content related to the query. Depending on the results of Z_t (and Y_{t-1}), the variable Y_t selects video shots to further summarize the remaining content in the video segment \mathcal{Y}_t . The arrows in each layer impose diversity by DPP between the shots selected from two adjacent video segments—we thus have Markov diversity, in contrast

to global diversity, in order to allow two (or more) visually similar shots to be simultaneously sampled to the summary if they appear at far-apart time steps (e.g., a man left home in the morning and returned home in the afternoon).

We define two types of DPPs for the two layers of SH-DPP, respectively.

Z-Layer to Summarize Query-Relevant Shots. We apply a conditional DPP $P(Z_t|q, Z_{t-1}, \mathcal{Y}_t)$ at each time step t over the ground set $\mathcal{Y}_t \cup \{Z_{t-1} = z_{t-1}\}$, where \mathcal{Y}_t consists of all the shots in partition t and z_{t-1} are the shots selected by Z_{t-1} . In other words, the DPP here is conditioned on the selected items z_{t-1} of the previous time step, enforcing Markov diversity between two consecutive time steps,

$$P(Z_t = z_t|q, Z_{t-1} = z_{t-1}, \mathcal{Y}_t) = \frac{\det \Omega_{z_{t-1} \cup z_t}}{\det(\Omega_{z_{t-1} \cup \mathcal{Y}_t} + I_t)} \quad (5)$$

where I_t is the same as an identity matrix except that its diagonal values are zeros at the entries indexed by z_{t-1} .

Different from seqDPP, we dedicate the Z-layer to query-relevant shots only. This is achieved by how we train SH-DPP (Sect. 3.3) and the way we parameterize the DPP kernel matrix,

$$\Omega_{ij} = [\mathbf{f}_i(q)]^T W^T W [\mathbf{f}_j(q)] \quad (6)$$

where $\mathbf{f}(q)$ is a query-dependent feature vector of a shot (Sect. 4.4). In testing, the Z-layer only selects shots that are relevant to the user query q , and leaves all the unselected shots to the Y-layer for further summarization.

Y-Layer to Summarize the Remaining Shots. The decision to include a shot in the query-focused video summarization is driven by two interplayed forces: the shot’s relevance to the query and its representativeness in the context. Given a user query q (e.g., CAR+FLOWER) and a long video \mathcal{Y} , likely many video shots are irrelevant to the query. As a result, we need another Y-layer to compensate the query-relevant shots selected by the Z-layer. In particular, we define the conditional probability distribution for the Y-layer variables as,

$$P(Y_t = y_t|Y_{t-1} = y_{t-1}, Z_t = z_t, \mathcal{Y}_t) = \frac{\det \Upsilon_{y_{t-1} \cup z_t \cup y_t}}{\det(\Upsilon_{y_{t-1} \cup \mathcal{Y}_t} + I'_t)} \quad (7)$$

where y_{t-1} is the selected subset in previous time step at the Y-layer, z_t is the selected subset of query-relevant shots in current time step by the Z-layer, and I'_t is a diagonal matrix with ones indexed by $\mathcal{Y}_t \setminus z_t$ and zeros everywhere else.

Conditioning the Y-layer on the Z-layer has two advantages. First, no redundant information that is already selected by Z-layer is added by the Y-layer again to the summary, i.e., the shots selected by Y-layer are diverse from those by Z-layer. Second, Y-layer can, to some extent, compensate the missed query-relevant shots by Z-layer that were supposed to be selected.

Note that the Y-layer involves a new DPP kernel Υ , different from that used for the Z-layer. The reason is twofold: first, two layers of variables serve to

select different (query relevant or important) types of shots, and second, the user may expect various levels of diversity from the summary. When a user searches for CAR+FLOWER, s/he probably would like to see more details in the shots of wedding car than in the shots of police, making it necessary to have two types of DPP kernels. The Y -layer kernel is parameterized by:

$$\mathbf{r}_{ij} = \mathbf{f}_i^T \mathbf{V}^T \mathbf{V} \mathbf{f}_j \quad (8)$$

and we will discuss how to extract features \mathbf{f} from a shot in Sect. 4.4.

3.3 Training and Testing SH-DPP

The training data in our experiments are in the form of $(q, \mathcal{Y}, z^q, y^q)$, where z^q and y^q respectively denote the query relevant and irrelevant shots in the summary. We learn the model parameters \mathbf{W} and \mathbf{V} of SH-DPP by maximum likelihood estimation (MLE):

$$\max_{\mathbf{W}, \mathbf{V}} \sum_q \sum_{\mathcal{Y}} \log P_{\text{SH}}(\{y_1, z_1\}, \dots, \{y_T, z_T\} | q, \mathcal{Y}) - \lambda_1 \|\mathbf{W}\|_F^2 - \lambda_2 \|\mathbf{V}\|_F^2, \quad (9)$$

where $\|\cdot\|_F^2$ is the squared Frobenius norm. We tune the hyper-parameters λ_1 and λ_2 by a leaving-one-video-out strategy, and optimize the above problem by gradient descent (cf. Supplement Material for more details on optimization).

After obtaining the local optimum \mathbf{W}^* and \mathbf{V}^* from the training, we need to know how to maximize the SH-DPP $P_{\text{SH}}(y|q, \mathcal{Y})$ for the testing stage (cf. Eq. (2)). However, the maximization remains a computationally extensive combinatorial problem. We thus follow [27] to have an approximate online inference procedure:

$$\begin{aligned} z_1^* &= \operatorname{argmax}_{z \in \mathcal{Y}_1} P(Z_1 = z | q, \mathcal{Y}_1), & y_1^* &= \operatorname{argmax}_{y \in \mathcal{Y}_1 \setminus z_1^*} P(Y_1 = y | z_1^*, \mathcal{Y}_1) \\ z_t^* &= \operatorname{argmax}_{z \in \mathcal{Y}_t} P(Z_t | q, z_{t-1}^*, \mathcal{Y}_t), & y_t^* &= \operatorname{argmax}_{y \in \mathcal{Y}_t \setminus z_t^*} P(Y_t | z_t^*, y_{t-1}^*, \mathcal{Y}_t), \quad t \geq 2 \end{aligned} \quad (10)$$

where we exhaustively search for z_t^* and y_t^* from \mathcal{Y}_t at each time step. Thanks to the online inference, SH-DPP can readily handle endlessly streaming videos.

4 Experiment Setup

In this section, we describe the datasets, features of a video shot, user queries, query-focused video summaries for training/evaluation, and finally, metrics to evaluate our learned video summarizer SH-DPP.

4.1 Datasets

We use the UT Egocentric (UTE) dataset [31] and TV episodes [32] whose dense user annotations are provided in [32]. The UTE dataset includes four daily life

egocentric videos, each 3–5 h long, and the TV episodes contain four videos, each roughly 45 min long. These two datasets are very different in nature. The videos in UTE are long and recorded in an uncontrolled environment from the first-person view. As a result, many of the visual scenes are repetitive and likely unwanted in the user summaries. In contrast, the TV videos are episodes of TV series from the third person’s viewpoint; the scenes are hence controlled and concise. A good summarizer should be able to work/learn well in both scenarios.

In [32], all the UTE/TV videos are partitioned to 5/10-second shots, respectively, and for each shot a textual description is provided by a human subject. Additionally, for each video, 3 reference summaries are also provided each as a subset of the textual annotations. Using dense text annotations, we are able to derive from the text both user queries and two types of query-focused video summaries, respectively, for patient and impatient users.

4.2 User Queries

In this paper, a user query comprises one or more noun concepts (e.g., CAR, FLOWER, KID); more generic queries are left for future research. There are many nouns in the text annotations of the video shots, but are they all useful for users to construct queries? Likely no. Any useful nouns have to be machine-detectable so that the system can “understand” the user queries. To this end, we construct a lexicon of concepts by overlapping all the nouns in the annotations with the nouns in SentiBank [41], which is a large collection of visual concepts and corresponding detectors. This results in a lexicon of 70/52 concepts for the UTE/TV dataset (see Table 1 in the supplementary material). Each pair of concepts is considered as a user query for both training and testing our SH-DPP video summarizer. Besides, at the testing phase, we also examine novel queries—all the triples of concepts.

4.3 Query-Focused Video Summaries

For each input query and video, we need to know the “groundtruth” video summary for training and evaluating SH-DPP. We construct such summaries based on the “oracle” summaries introduced in [27].

Oracle Summary. As mentioned earlier in Sect. 4.1, there are three human-annotated summaries y^u , $u = 1, 2, 3$ for each video \mathcal{Y} . An “oracle” summary y^o has the maximum agreement with all of the three annotated summaries, and can be understood as the summary by an “average” user. Such a summary is found by a greedy algorithm [38]. Initialize $y^o = \emptyset$. In each iteration, the set y^o increases by one video shot i which gives rise to the largest marginal gain $G(i)$,

$$y^o \leftarrow y^o \cup \underset{i \in \mathcal{Y}}{\operatorname{argmax}} G(i), \quad G(i) = \sum_u \operatorname{F-score}(y^o \cup i, y_u) - \operatorname{F-score}(y^o, y_u) \quad (11)$$

where the F-score follows [32] and is explained in Sect. 4.5. The algorithm stops when there is no such shot that the gain $G(i)$ is greater than 0. Note that thus far the oracle summary is independent of the user query.

Query-Focused Video Summary. We consider two types of users. A **patient user** would like to watch all the shots relevant to the query in addition to the summary of the other visual content of the video. For example, all the shots whose textual descriptions have the word CAR should be included in the summary if CAR shows up in the query. We union such shots with the oracle summary to have the query-focused summary for the patient user. On the other extreme, an **impatient user** may only want to check the existence of the relevant shots, in contrast to watching all of them. To conduct experiments for the impatient users, we overlap the concepts in the oracle summary with the concept lexicon (cf. Sect. 4.2), and generate all possible bi-concept queries from the survived concepts. Note that the oracle summaries are thus the gold standards for training video summarizers for the impatient users.

4.4 Features

We extract high-level concept-oriented features \mathbf{h} and contextual features \mathbf{l} for a video shot. For each concept in the lexicon (of size 70/52 for the UTE/TV dataset), we firstly use its corresponding SentiBank detector(s) [41] to obtain the detection scores of the key frames, and then average them within each shot. Some of the concepts each maps to more than one detectors. For instance, there are beautiful SKY, clear SKY, and sunny SKY detectors for the concept SKY. We max-pool their shot-level scores, so there is always one detection score, which is between 0 and 1, for each concept. The resultant high-level concept-oriented feature vector \mathbf{h} is 70D/52D for a shot of a UTE/TV video. We ℓ_2 normalize it.

Furthermore, we design some contextual features \mathbf{l} for a video shot based on the low-level features that SentiBank uses as input to its classifiers. This set of low-level features includes color histogram, GIST [42], LBP [43], Bag-of-Words descriptor, and an attribute feature [44]. With these features, we put a temporal window around each frame, and compute the mean-correlation as a contextual feature for the frame. The mean-correlation shows how well the frame is representative of the other frames in the temporal window. By varying the window size from 5 to 15 with step size 2, we obtain a 6D feature vector. Again we average pool them within each shot, followed by ℓ_2 normalization, to have the shot-level contextual feature vector \mathbf{l} .

The concept-oriented and contextual features are concatenated as the overall shot-level feature vector $\mathbf{f} \equiv [\mathbf{h}; \mathbf{l}]$ for parameterizing the DPP kernel of the Y -layer (Eq. (8)). The Z -layer kernel calls for query-dependent features $\mathbf{f}(q)$ (Eq. (6)). For this purpose, we scale the concept-oriented features according to the query: $\mathbf{f}(q) \equiv \mathbf{h} \circ \boldsymbol{\alpha}(q)$, where \circ is the element-wise product between two vectors, and the scaling factors $\boldsymbol{\alpha}(q)$ are 1 for the concepts shown in the query and 0.5 otherwise (see Fig. 1(a, b) for an example). Though we may employ more sophisticated query-dependent features, the simple features scaled by the query perform well in our experiments. The simplicity also enables us to feed the same features to vanilla and sequential DPPs for fair comparison.

4.5 Evaluation

We evaluate a system generated video summary by contrasting it against the “groundtruth” summary. The comparison is based on the dense text annotations [32]. In particular, the video summaries are mapped to text paragraphs and then compared by the ROUGE-SU metric [45]. We report the precision, recall, F-score returned by ROUGE-SU.

In addition, we also introduce a new metric, called *hitting recall*, to evaluate the system summaries from the query-focused perspective. Given the input query q and long video \mathcal{V} , denote by S^q the shots relevant to the query in the “groundtruth” summary, and S_{SYSTEM}^q the query-relevant shots hit by a video summarizer. The hitting recall is calculated by $\text{HR} = |S_{\text{SYSTEM}}^q|/|S^q|$, where $|\cdot|$ is the cardinality of a set. For our SH-DPP model, we report the hitting recalls for both the overall summaries and those by the Z -layer only.

Table 1. Results of query-focused video summarization with **bi-concept** queries.

Patient users	UTE (%)					TV episodes (%)				
	F	Prec	Recall	HR	HR _Z	F	Prec	Recall	HR	HR _Z
Sampling	22.12	35.07	17.11	23.61	n/a	27.99	34.75	24.36	16.00	n/a
Ranking	20.66	24.35	18.38	22.05	n/a	32.19	39.96	32.19	16.61	n/a
SubMod [28]	20.98	31.40	26.99	30.10	n/a	32.19	41.59	27.01	21.69	n/a
Quasi [46]	12.45	19.47	13.14	14.95	n/a	31.88	27.49	41.69	19.67	n/a
DPP [38]	15.7	19.22	32.08	30.94	n/a	29.62	35.26	34.00	21.29	n/a
seqDPP [27]	18.85	20.59	35.83	31.91	n/a	27.96	23.80	35.62	14.08	n/a
SH-DPP (ours)	21.27	17.87	41.65	38.26	36.92	37.02	38.41	36.82	23.76	20.35
Impatient users	UTE (%)					TV episodes (%)				
	F	P	R	HR	HR _Z	F	P	R	HR	HR _Z
Sampling	25.44	44.16	18	6.48	n/a	33.74	41.03	28.8	13.03	n/a
Ranking	17.92	21.86	15.46	4.4	n/a	29.67	37.56	24.72	15.43	n/a
SubMod [28]	27.10	51.79	18.85	8.05	n/a	29.41	38.51	23.85	8.65	n/a
Quasi [46]	11.52	42.32	7.06	1.63	n/a	25.09	27.25	23.71	17.06	n/a
DPP [38]	14.36	30.9	16.18	12.54	n/a	26.01	28.85	39.15	18.86	n/a
seqDPP [27]	12.93	7.89	43.39	12.68	n/a	23.35	16.60	39.69	12.56	n/a
SH-DPP (ours)	25.56	18.51	45.21	22.91	11.57	35.36	30.94	42.02	17.07	17.07

4.6 Implementation Details

Here we report some details in our implementation of SH-DPP. Out of the four videos in either UTE or TV, we use three videos for training and the remaining one for testing. Each video is taken for testing once and then the averaged results are reported. In the training phase, there are two hyper-parameters in our approach: λ_1 and λ_2 (cf. Eq. (9)). We choose their values by the leave-one-video-out strategy (over the training videos only). The lower-dimensions of \mathbf{W}

and \mathbf{V} are both fixed to 10, the same as used in seqDPP [27]. Varying this number has little effect to the results and we leave the study about it to the Supplementary Material. We put 10 shots in the ground set \mathcal{Y}_t at each time step, and also examine the ground sets of the other sizes in the experiments.

We train our model SH-DPP using bi-concept queries. However, we test it using not only the bi-concept queries but also novel three-concept queries.

5 Experimental Results

This section presents the comparison results of our approach and some competitive baselines, effect of the ground set size, and finally qualitative results.

5.1 Comparison Results

Table 1 shows the results of different summarizers for the query-focused video summarization when the patient and impatient users supply **bi-concept** queries, while Table 2 includes the results for novel **three-concept** queries. Note that only bi-concept queries are used to train the summarizers. We report the results on both UTE and TV datasets, and contrast our SH-DPP to the following methods: (1) uniformly sampling K shots, (2) ranking, where for each query we apply the corresponding concept detectors to the shots, assign to a shot a ranking score as the maximum detection score, and then keep the top K shots, (3) vanilla DPP [38], where we remove the dependency between adjacent subset selection variables in Fig. 2(a), (4) seqDPP [27], (5) SubMod [28], where convex combination of a set of objectives is learned from user summaries, and (6) Quasi [46] which is an unsupervised method based on group sparse coding. We let K be the number of shots in the groundtruth summary; therefore, such privileged information makes (1), (2), and (5) actually strong baselines. We use the same ground sets, whose sizes are fixed to 10 for DPP, seqDPP, and our SH-DPP. All the results are evaluated by the F-score, Precision, and Recall of ROUGE-SU, as well as the hitting recall (HR) (cf. Sect. 4.5).

Interesting insights can be inferred from Tables 1 and 2. An immediate observation is that our SH-DPP is able to generate better overall summaries as our average F-scores are higher than the others'. Furthermore, our method is able to adapt itself to two essentially different datasets, the UTE daily life egocentric videos and TV episodes.

On UTE, we expect both SH-DPP and seqDPP to outperform vanilla DPP, because the egocentric videos are very long and include many unwanted scenes,

Table 2. Results of query-focused video summarization with novel **three-concept** queries.

Patient users	UTE (%)			TV episodes (%)		
	F	HR	HR _Z	F	HR	HR _Z
DPP	20.7	38.53	n/a	29.84	21.68	n/a
seqDPP	18.03	30.3	n/a	24.29	14.15	n/a
SH-DPP (ours)	24.54	41.23	40.43	36.3	24.73	21.71
Impatient users	UTE (%)			TV episodes (%)		
	F	HR	HR _Z	F	HR	HR _Z
DPP	14.77	17.28	n/a	24.71	18.31	n/a
seqDPP	19.4	19.17	n/a	29.31	10.09	n/a
SH-DPP (ours)	29.59	25.82	15.36	33.94	17.39	12.33

making the dependency between different subset selection variables essential for eliminating repetitions. In contrast, as mentioned in Sect. 4.1, the TV episodes are from the world of professional recording, and the scenes rapidly change from shots to shots. Therefore, in this case, the dependency is weak and DPP may be able to catch up seqDPP’s performance. These hypotheses are verified in the results, if we compare DPP with seqDPP in Tables 1 and 2.

Another important observation is that in 6 out of the 8 experiments: {patient and impatient users} on {UTE and TV datasets} by {bi-concept and novel three-concept queries}, the proposed SH-DPP has better hitting recalls than the other methods, indicating a better response to the user queries. Moreover, the hitting recalls are mainly captured by the Z-layer—the columns HR_Z are the hitting recalls of the shots selected by the Z-layer only of SH-DPP.

As it can be noticed from Table 1, uniform sampling has competitive performance compared to the other baselines and even outperforms SH-DPP in one scenario. The relatively good performance of random sampling can be explained by looking into the evaluation metric. ROUGE essentially evaluates the summaries by common word/phrase count and penalizing long or short summaries. Thus, accessing the number of groundtruth shots gives an advantage to Sampling. Figure 3 illustrates the change of performance when we deviate from the number of shots in groundtruth summary. This figure was generated using the TV Episodes dataset for both patient and impatient user cases.




Fig. 3. The effect of number of selected shots on the performance of uniform sampling.

5.2 A Peek into the SH-DPP Summarizer

Figure 4 is an exemplar summary for the query **FLOWERS+WALL** by SH-DPP. For each shot in the summary, we show the middle frame of that shot and the corresponding textual description. The groundtruth summary is also included at the bottom half of the figure. We can see that some query-relevant shots are successfully selected by the Z-layer. Conditioning on those, the Y-layer summarizes the remaining video. We highlight the text descriptions (in the blue color) that have exact matches in the groundtruth. However, please note that the other sentences are also highly correlated with some groundtruth sentences, for instance, “I looked at flowers at the booth” selected by the Z-layer versus “my friend and I looked at flowers at the booth” in the groundtruth summary.

One may wonder why the top-right shot is selected by the Z-layer, since it is visually not relevant to either **FLOWERS** or **WALL**. Inspection tells that it is due to the failure of the concept detectors; the detection scores are 0.86 and 0.65 out of 1 for **FLOWERS** and **WALL**, respectively. We may improve our SH-DPP for the query-focused video summarization by using better concept detectors.

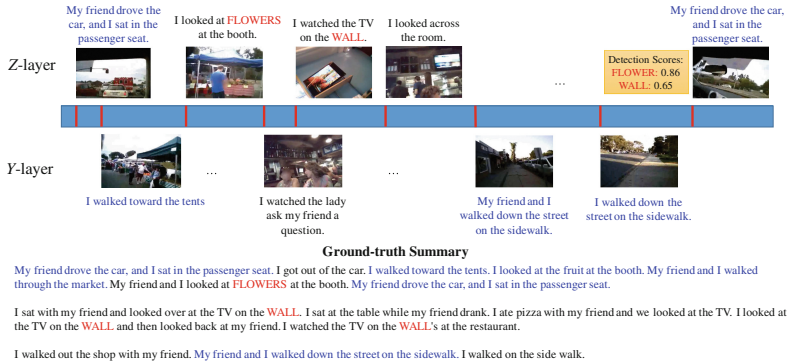


Fig. 4. A peek into SH-DPP. Given the query **FLOWERS+WALL**, the Z-layer of SH-DPP is supposed to summarize the shots relevant to the query. Conditioning on those results, the Y-layer summarizes the remaining video. (Color figure online)

In the Supplementary Material: We first show all the concepts used in our experiments. Also, we describe the detailed training algorithm for SH-DPP using gradient descent. The resultant optimization problem is non-convex; we explain how to choose the initializations by the leaving-one-video-out strategy. We also show that the SH-DPP results remain stable for different lower-dimensions of \mathbf{W} and \mathbf{V} . Furthermore, we show how changing the groundset size affects the performance of SH-DPP. Finally, more qualitative results are included in the Supplementary Material.

6 Conclusions

In this paper, we examined a query-focused video summarization problem, in which the decision to select a video shot to the summary depends on both (1) the relevance between the shot and the query and (2) the importance of the shot in the context of the video. To tackle this problem, we developed a probabilistic model, Sequential and Hierarchical Determinantal Point Process (SH-DPP), as well as efficient learning and inference algorithms for it. Our SH-DPP summarizer can conveniently handle extremely long videos or online streaming videos. On two benchmark datasets for video summarization, our approach significantly outperforms some competing baselines. To the best of our knowledge, ours is the first work on query-focused video summarization, and has a great potential to be used in search engines, e.g., to display snippets of videos.

Acknowledgements. A.S. & B.G. are supported in part by NSF IIS #1566511. M.S. is partially supported by NIJ W911NF-14-1-0294.

References

1. Pritch, Y., Rav-Acha, A., Gutman, A., Peleg, S.: Webcam synopsis: peeking around the world. In: IEEE 11th International Conference on Computer Vision 2007, ICCV 2007, pp. 1–8. IEEE (2007)
2. Pal, C., Jojic, N.: Interactive montages of sprites for indexing and summarizing security video. In: IEEE Computer Society Conference on CVPR 2005, vol. 2. IEEE (2005)
3. Kang, H.W., Matsushita, Y., Tang, X., Chen, X.Q.: Space-time video montage. In: IEEE Computer Society Conference on CVPR 2006, vol. 2. IEEE (2006)
4. Jiang, R.M., Sadka, A.H., Crookes, D.: Advances in video summarization and skimming. In: Grgic, M., Delac, K., Ghanbari, M. (eds.) Recent Advances in Multimedia Signal Processing and Communications. SCI, vol. 231, pp. 27–50. Springer, Heidelberg (2009)
5. Rav-Acha, A., Pritch, Y., Peleg, S.: Making a long video short: dynamic video synopsis. In: 2006 IEEE Computer Society Conference on CVPR, vol. 1. IEEE (2006)
6. Goldman, D.B., Curless, B., Salesin, D., Seitz, S.M.: Schematic storyboarding for video visualization and editing. *ACM Trans. Graph. (TOG)* **25**, 862–871 (2006). ACM
7. Liu, T., Kender, J.R.: Optimization algorithms for the selection of key frame sequences of variable length. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2353, pp. 403–417. Springer, Heidelberg (2002)
8. Auer, A., Kender, J.R.: Video summaries through mosaic-based shot and scene clustering. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2353, pp. 388–402. Springer, Heidelberg (2002)
9. Vasconcelos, N., Lippman, A.: A spatiotemporal motion model for video summarization. In: Proceedings of IEEE Computer Society Conference on CVPR 1998, pp. 361–366. IEEE (1998)
10. Wolf, W.: Key frame selection by motion analysis. In: Proceedings of 1996 IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 1996, vol. 2, pp. 1228–1231. IEEE (1996)
11. Lee, K.M., Kwon, J.: A unified framework for event summarization and rare event detection. In: 2012 IEEE Conference on CVPR. IEEE (2012)
12. Cong, Y., Yuan, J., Luo, J.: Towards scalable summarization of consumer videos via sparse dictionary selection. *IEEE Trans. Multimedia* **14**(1), 66–75 (2012)
13. Ngo, C., Ma, Y., Zhang, H.: Automatic video summarization by graph modeling. In: Proceedings of the Ninth IEEE International Conference on Computer Vision 2003. IEEE (2003)
14. Khosla, A., Hamid, R., Lin, C.J., Sundaresan, N.: Large-scale video summarization using web-image priors. In: Proceedings of the IEEE Conference on CVPR (2013)
15. Kim, G., Sigal, L., Xing, E.: Joint summarization of large-scale collections of web images and videos for storyline reconstruction. In: Proceedings of the IEEE Conference on CVPR (2014)
16. Xiong, B., Grauman, K.: Detecting snap points in egocentric video with a web photo prior. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part V. LNCS, vol. 8693, pp. 282–298. Springer, Heidelberg (2014)
17. Chu, W.S., Song, Y., Jaimes, A.: Video co-summarization: video summarization by visual co-occurrence. In: Proceedings of the IEEE Conference on CVPR (2015)

18. Song, Y., Vallmitjana, J., Stent, A., Jaimes, A.: TVSum: summarizing web videos using titles. In: Proceedings of the IEEE Conference on CVPR (2015)
19. Liu, W., Mei, T., Zhang, Y., Che, C., Luo, J.: Multi-task deep visual-semantic embedding for video thumbnail selection. In: Proceedings of the IEEE Conference on CVPR (2015)
20. Potapov, D., Douze, M., Harchaoui, Z., Schmid, C.: Category-specific video summarization. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part VI. LNCS, vol. 8694, pp. 540–555. Springer, Heidelberg (2014)
21. Sun, M., Farhadi, A., Seitz, S.: Ranking domain-specific highlights by analyzing edited videos. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part I. LNCS, vol. 8689, pp. 787–802. Springer, Heidelberg (2014)
22. Xu, J., Mukherjee, L., Li, Y., Warner, J., Rehg, J.M., Singh, V.: Gaze-enabled egocentric video summarization via constrained submodular maximization. In: Proceedings of the IEEE Conference on CVPR (2015)
23. Gygli, M., Grabner, H., Riemenschneider, H., Van Gool, L.: Creating summaries from user videos. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part VII. LNCS, vol. 8695, pp. 505–520. Springer, Heidelberg (2014)
24. Lu, Z., Grauman, K.: Story-driven summarization for egocentric video. In: Proceedings of the IEEE Conference on CVPR (2013)
25. Lee, Y.J., Grauman, K.: Predicting important objects for egocentric video summarization. *Int. J. Comput. Vis.* **114**(1), 38–55 (2015)
26. Liu, D., Hua, G., Chen, T.: A hierarchical visual model for video object summarization. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(12), 2178–2190 (2010)
27. Gong, B., Chao, W.L., Grauman, K., Sha, F.: Diverse sequential subset selection for supervised video summarization. In: Advances in Neural Information Processing Systems, pp. 2069–2077 (2014)
28. Gygli, M., Grabner, H., Van Gool, L.: Video summarization by learning submodular mixtures of objectives. In: Proceedings of the IEEE Conference on CVPR (2015)
29. Nenkova, A., McKeown, K.: A survey of text summarization techniques. In: Aggarwal, C.C., Zhai, C.X. (eds.) *Mining Text Data*, pp. 43–76. Springer, Heidelberg (2012)
30. Kulesza, A., Taskar, B.: Determinantal point processes for machine learning. arXiv preprint [arXiv:1207.6083](https://arxiv.org/abs/1207.6083) (2012)
31. Ghosh, J., Lee, Y.J., Grauman, K.: Discovering important people and objects for egocentric video summarization. In: 2012 IEEE Conference on CVPR. IEEE (2012)
32. Yeung, S., Fathi, A., Fei-Fei, L.: Videoseq: Video summary evaluation through text. arXiv preprint [arXiv:1406.5824](https://arxiv.org/abs/1406.5824) (2014)
33. Daumé III, H., Marcu, D.: Bayesian query-focused summarization. In: Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics (2006)
34. Schilder, F., Kondadadi, R.: Fastsum: fast and accurate query-based multi-document summarization. In: Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers, pp. 205–208. Association for Computational Linguistics (2008)
35. Gupta, S., Nenkova, A., Jurafsky, D.: Measuring importance and query relevance in topic-focused multi-document summarization. In: Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions. Association for Computational Linguistics, pp. 193–196 (2007)

36. Ellouze, M., Boujemaa, N., Alimi, A.M.: IM(S)²: interactive movie summarization system. *J. Vis. Commun. Image Represent.* **21**(4), 283–294 (2010)
37. Xiong, B., Kim, G., Sigal, L.: Storyline representation of egocentric videos with an applications to story-based search. In: *Proceedings of the IEEE International CVPR* (2015)
38. Kulesza, A., Taskar, B.: Learning determinantal point processes. arXiv preprint [arXiv:1202.3738](https://arxiv.org/abs/1202.3738) (2012)
39. Chao, W.L., Gong, B., Grauman, K., Sha, F.: Large-margin determinantal point processes. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)* (2015)
40. Affandi, R.H., Kulesza, A., Fox, E.B.: Markov determinantal point processes. arXiv preprint [arXiv:1210.4850](https://arxiv.org/abs/1210.4850) (2012)
41. Borth, D., Chen, T., Ji, R., Chang, S.F.: Sentibank: large-scale ontology and classifiers for detecting sentiment and emotions in visual content. In: *Proceedings of the 21st ACM International Conference on Multimedia*. ACM (2013)
42. Oliva, A., Torralba, A.: Modeling the shape of the scene: a holistic representation of the spatial envelope. *Int. J. Comput. Vision* **42**(3), 145–175 (2001)
43. Ojala, T., Pietikäinen, M., Mäenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(7), 971–987 (2002)
44. Yu, F., Cao, L., Feris, R., Smith, J., Chang, S.F.: Designing category-level attributes for discriminative visual recognition. In: *Proceedings of the IEEE Conference on CVPR* (2013)
45. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: *Proceedings of the ACL-04 Workshop, Text Summarization Branches Out*, vol. 8 (2004)
46. Zhao, B., Xing, E.: Quasi real-time summarization for consumer videos. In: *Proceedings of the IEEE Conference on CVPR* (2014)